


Building Robust Datasets for Commonsense Reasoning

Rowan Zellers

rowanzellers.com





A Venn diagram illustrating the relationship between 'The world', 'Language', and 'Vision'. A large yellow circle represents 'The world'. Inside it, two smaller circles overlap: a green one on the left labeled 'Language' and an orange one on the right labeled 'Vision'. The intersection of the green and orange circles is a darker brownish-green color.

The world

Language

Vision

A man has a few tools and is pumping his car up so he can take off the tire.
He...



A man has a few tools and is pumping his car up so he can take off the tire.
He...



- A. stops on the front of the bike and moves it to the left.
- B. gets out of the car while leaving the engine running.
- C. uses the tool to take off all of the nuts one by one.
- D. goes down from the cars, landing straight in.



A man has a few tools and is pumping his car up so he can take off the tire.
He...

- A. stops on the front of the bike and moves it to the left.
- B. gets out of the car while leaving the engine running.
- C. uses the tool to take off all of the nuts one by one.
- D. goes down from the cars, landing straight in.



A man has a few tools and is pumping his car up so he can take off the tire.
He...



A. stops on the front of the bike and



This is natural language inference that requires commonsense reasoning!

one by one.

D. goes down from the cars, landing straight in.



SWAG: A Large-Scale Adversarial
Dataset for Grounded Commonsense
Inference (emnlp18)

Our contributions with *SWAG*

- SWAG: Natural Language Inference + Commonsense Reasoning



Our contributions with *SWAG*

- SWAG: Natural Language Inference + Commonsense Reasoning
- Adversarial Filtering




Scope of Natural Language Inference (NLI)

Much of today's NLI requires only linguistic knowledge, without as much commonsense reasoning.


Scope of Natural Language Inference (NLI)

Much of today's NLI requires only linguistic knowledge, without as much commonsense reasoning. **Examples from SNLI:**



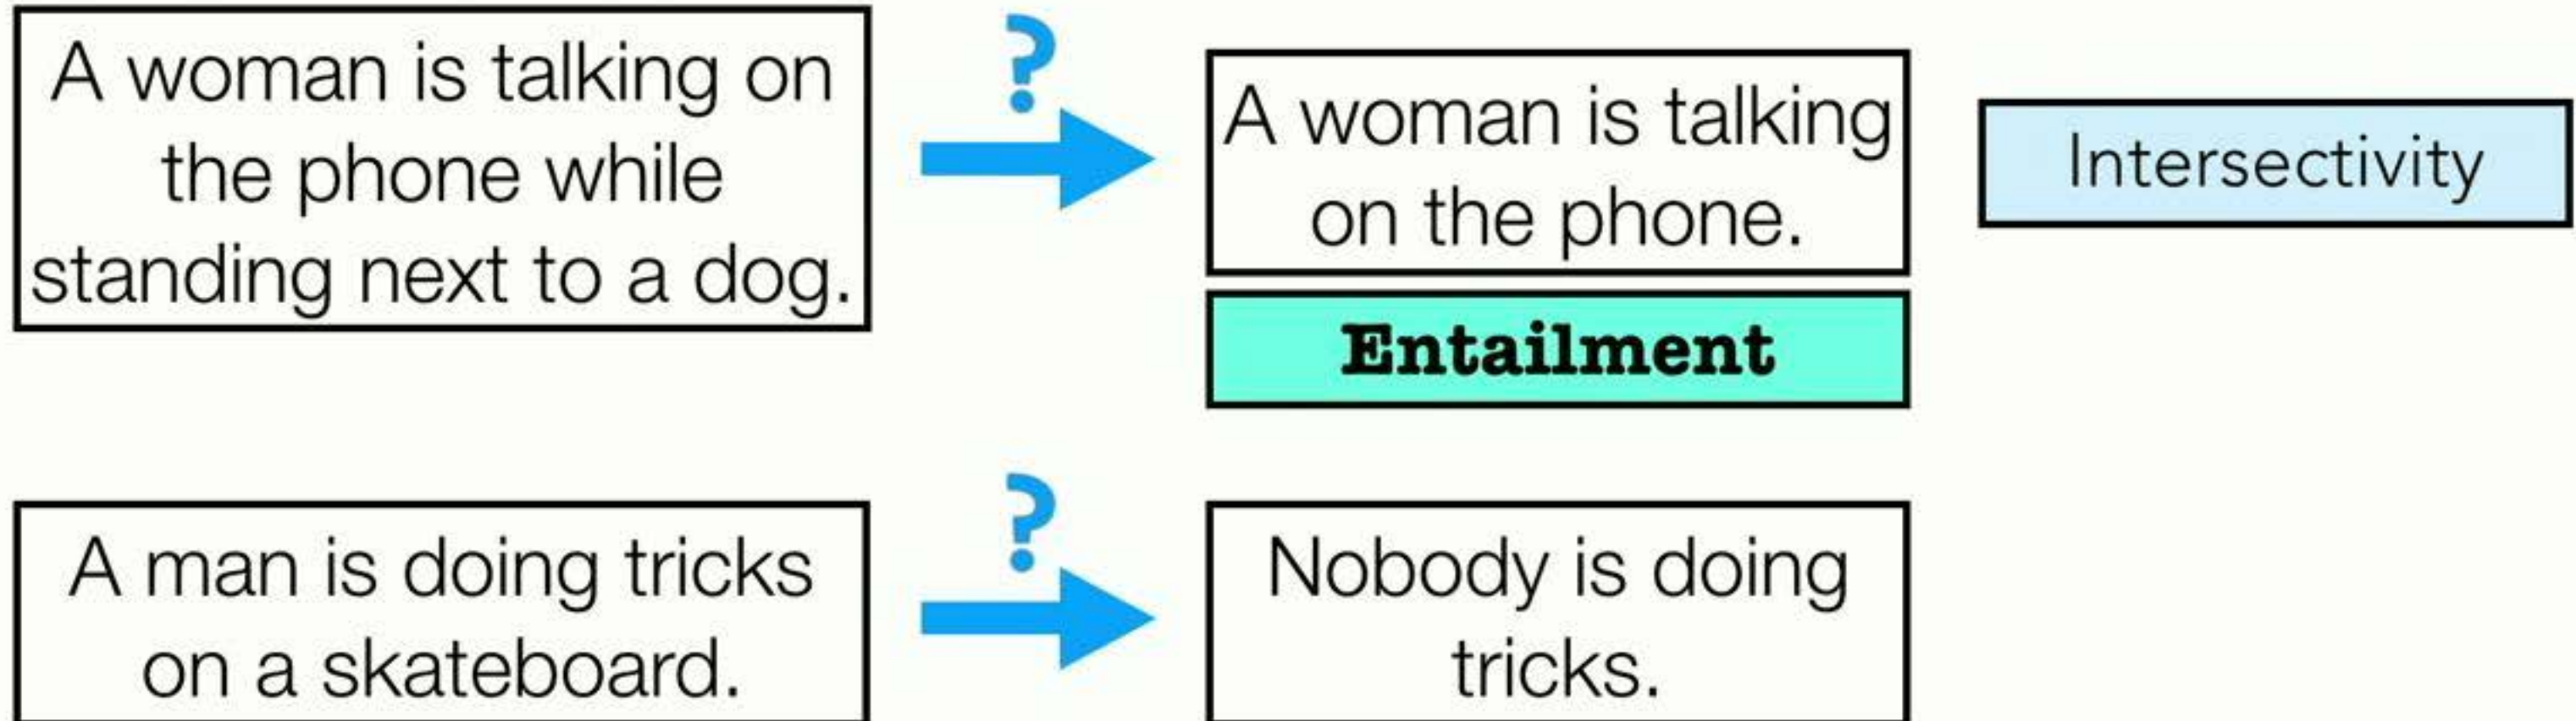
Scope of Natural Language Inference (NLI)

Much of today's NLI requires only linguistic knowledge, without as much commonsense reasoning. *Examples from SNLI:*



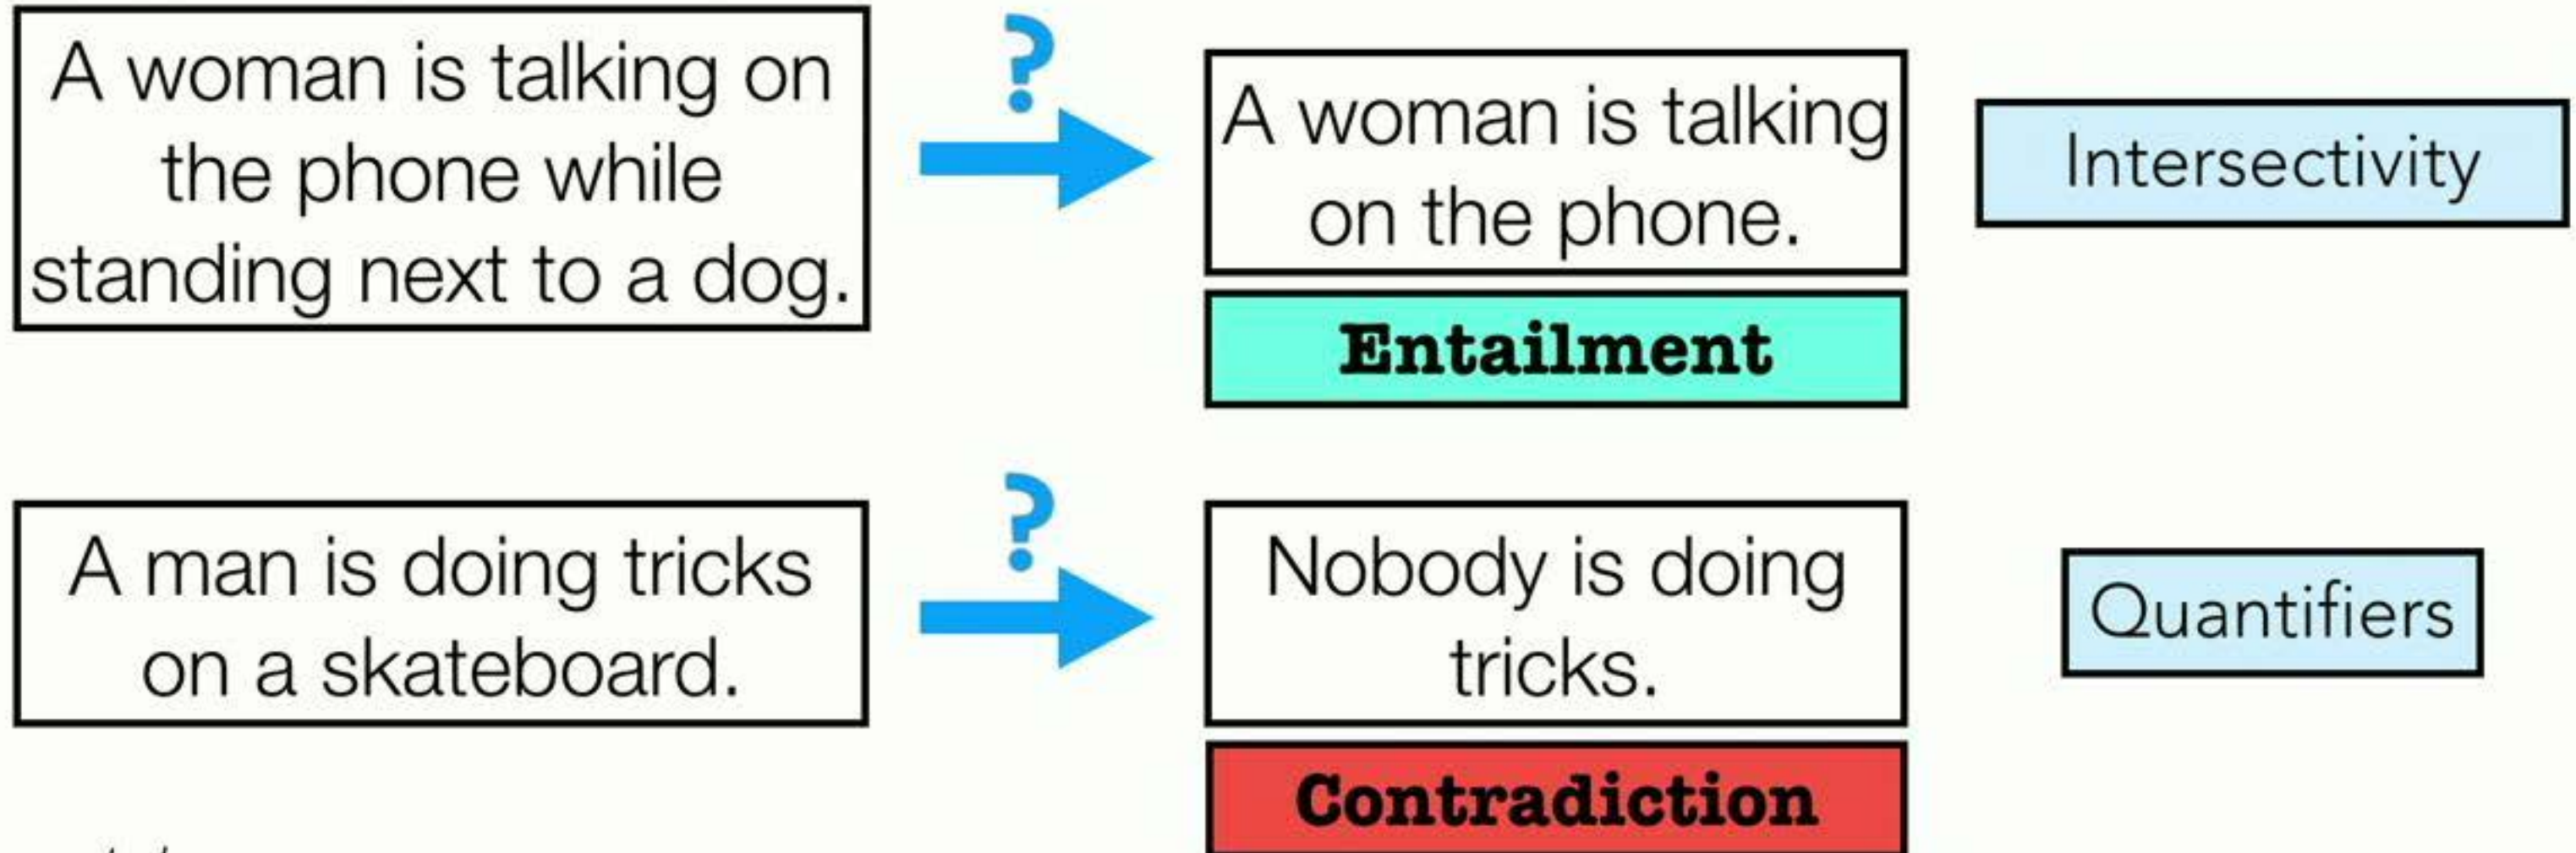
Scope of Natural Language Inference (NLI)

Much of today's NLI requires only linguistic knowledge, without as much commonsense reasoning. *Examples from SNLI:*




Scope of Natural Language Inference (NLI)

Much of today's NLI requires only linguistic knowledge, without as much commonsense reasoning. *Examples from SNLI:*



Re-emphasizing Commonsense in NLI

Natural Language Inference
was supposed to be:




(Dagan et. al, 2006; LoBue and Yates, 2011)



Situations with Adversarial Generations (*SWAG*)

A man has a few tools and is pumping his car up so he can take off the tire. He



t



ACTIVITYNET

LSMDC

Krishna et al., 2017,
Rohrbach et al., 2016

Situations with Adversarial Generations (*SWAG*)

A man has a few tools and is pumping his car up so he can take off the tire. He

- A. stops on the front of the bike and moves it to the left.
- B. gets out of the car while leaving the engine running.
- C. uses the tool to take off all of the nuts one by one.
- D. goes down from the cars, landing straight in




Situations with Adversarial Generations (**SWAG**)

A man has a few tools and is pumping his car up so he can take off the tire. He



- A. stops on the front of the bike and moves it to the left.
- B. gets out of the car while leaving the engine running.

C. uses the tool to take off all of the nuts one by one.



- D. goes down from the cars, landing straight in

SWAG is purely a natural language inference task: the video is never used.



Situations with Adversarial Generations (**SWAG**)

A man has a few tools and is pumping his car up so he can take off the tire. He



A. stops on the front of the bike and moves it to the left

B.

How do we get the wrong answers?

C.

nuts one by one.

D.

goes down from the cars, landing straight in



Our contributions

- SWAG: Natural Language Inference + Commonsense Reasoning
- Adversarial Filtering



Annotation Artifacts

Human written datasets are susceptible to *annotation artifacts*: stylistic patterns that give unwanted clues for the labels.

Schwartz et al., 2017



Gururangan et al., 2018

Poliak et al., 2018

Annotation Artifacts

Human written datasets are susceptible to *annotation artifacts*: stylistic patterns that give unwanted clues for the labels.


Schwartz et al., 2017
Gururangan et al., 2018
Poliak et al., 2018




Annotation Artifacts

Human written datasets are susceptible to *annotation artifacts*: stylistic patterns that give unwanted clues for the labels.

Schwartz et al., 2017
Gururangan et al., 2018
Poliak et al., 2018



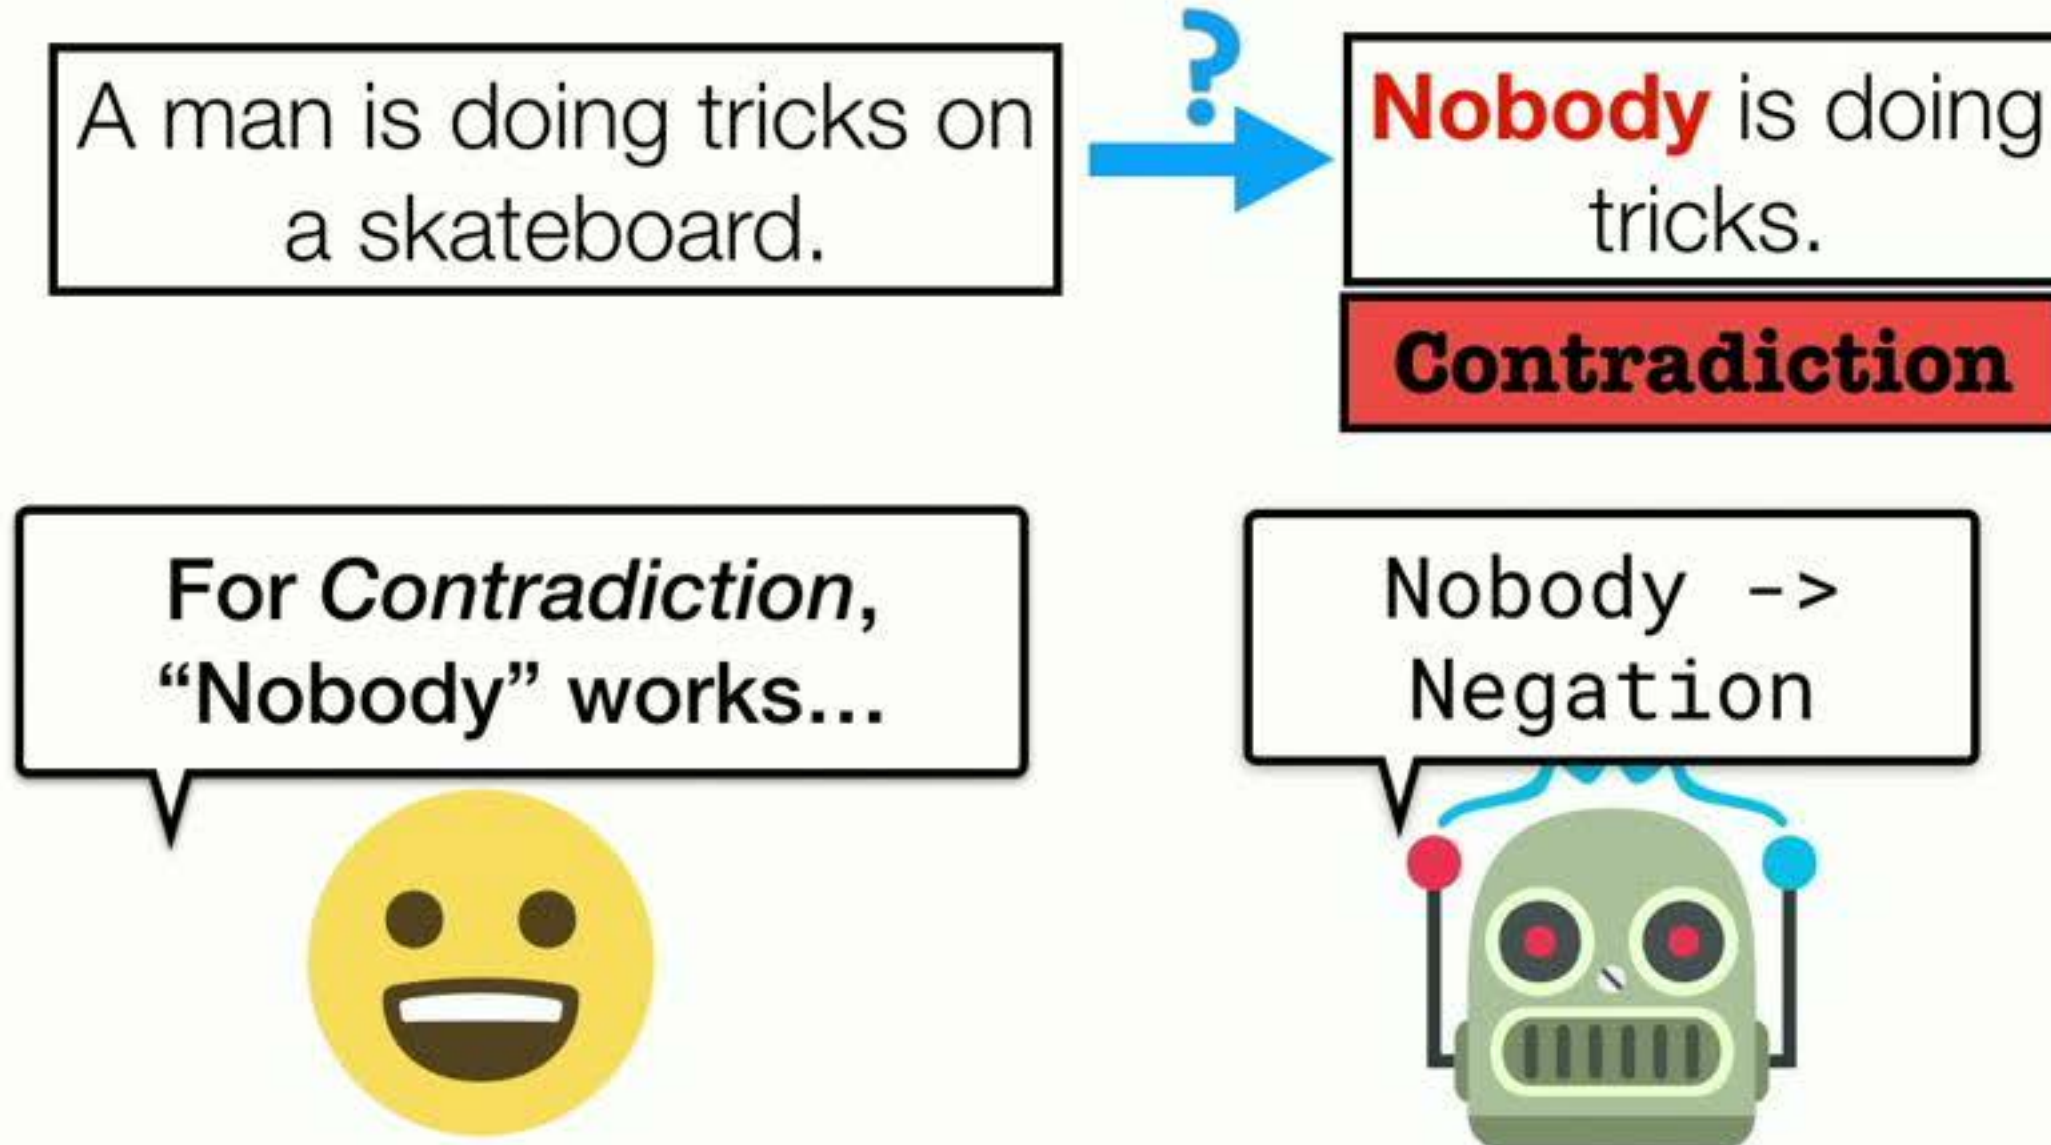
For *Contradiction*,
"Nobody" works...




Annotation Artifacts

Human written datasets are susceptible to *annotation artifacts*: stylistic patterns that give unwanted clues for the labels.


Schwartz et al., 2017
Gururangan et al., 2018
Poliak et al., 2018



A different approach...




A different approach...



Let's have machines
write the wrong
endings!



A different approach...



Let's have machines
write the wrong
endings!



Humans will **verify** that the
original endings are better
than the wrong ones.

A different approach...




Let's have machines
write the wrong
endings!

Machine written text still has artifacts!




Humans will **verify** that the
original endings are better
than the wrong ones.

Adversarial Filtering




We'll *massively oversample* candidate endings.

Adversarial Filtering




We'll *massively oversample* candidate endings.




And we'll use more models to remove obvious artifacts!

Adversarial Filtering



We'll *massively oversample* candidate endings.



And we'll use more models to remove obvious artifacts!

Adversarial Filtering

A man has a few tools and is pumping his car up so he can take off the tire.

He

NP

uses the tool to take off all of the nuts one by one.

Ground truth VP



Adversarial Filtering

A man has a few tools and is pumping his car up so he can take off the tire.

He

NP

uses the tool to take off all of the nuts one by one.

Ground truth VP



**LM
generated
endings**

- goes out onto the street.
- sits looking at the man.
- goes down from the cars, landing straight in.
- hauls the car over the bridge.

....

Adversarial Filtering

**Ground
truth VP**

I

uses the tool to take off all of the nuts one by one.



x_1^+

**LM
generated
endings**

- goes out onto the street.
- sits looking at the man.
- goes down from the cars, landing straight in.
- hauls the car over the bridge.

....

$x_{1,1}^-$

$x_{1,2}^-$

$x_{1,3}^-$

$x_{1,4}^-$

⋮

Adversarial Filtering

*Ground
truth VP*

I

$$x_1^+$$

$$x_{1,1}^-$$

$$x_{1,2}^-$$

$$x_{1,3}^-$$

$$x_{1,4}^-$$

⋮

*LM
generated
endings*

**Ground
truth VPs**

I

x_1^+

x_2^+

x_3^+

x_4^+

x_5^+

x_6^+

• • •

$x_{1,1}^-$

$x_{2,1}^-$

$x_{3,1}^-$

$x_{4,1}^-$

$x_{5,1}^-$

$x_{6,1}^-$

• • •

$x_{1,2}^-$

$x_{2,2}^-$

$x_{3,2}^-$

$x_{4,2}^-$

$x_{5,2}^-$

$x_{6,2}^-$

• • •

$x_{1,3}^-$

$x_{2,3}^-$

$x_{3,3}^-$

$x_{4,3}^-$

$x_{5,3}^-$

$x_{6,3}^-$

• • •

$x_{1,4}^-$

$x_{2,4}^-$

$x_{3,4}^-$

$x_{4,4}^-$

$x_{5,4}^-$

$x_{6,4}^-$

• • •

⋮

⋮

⋮

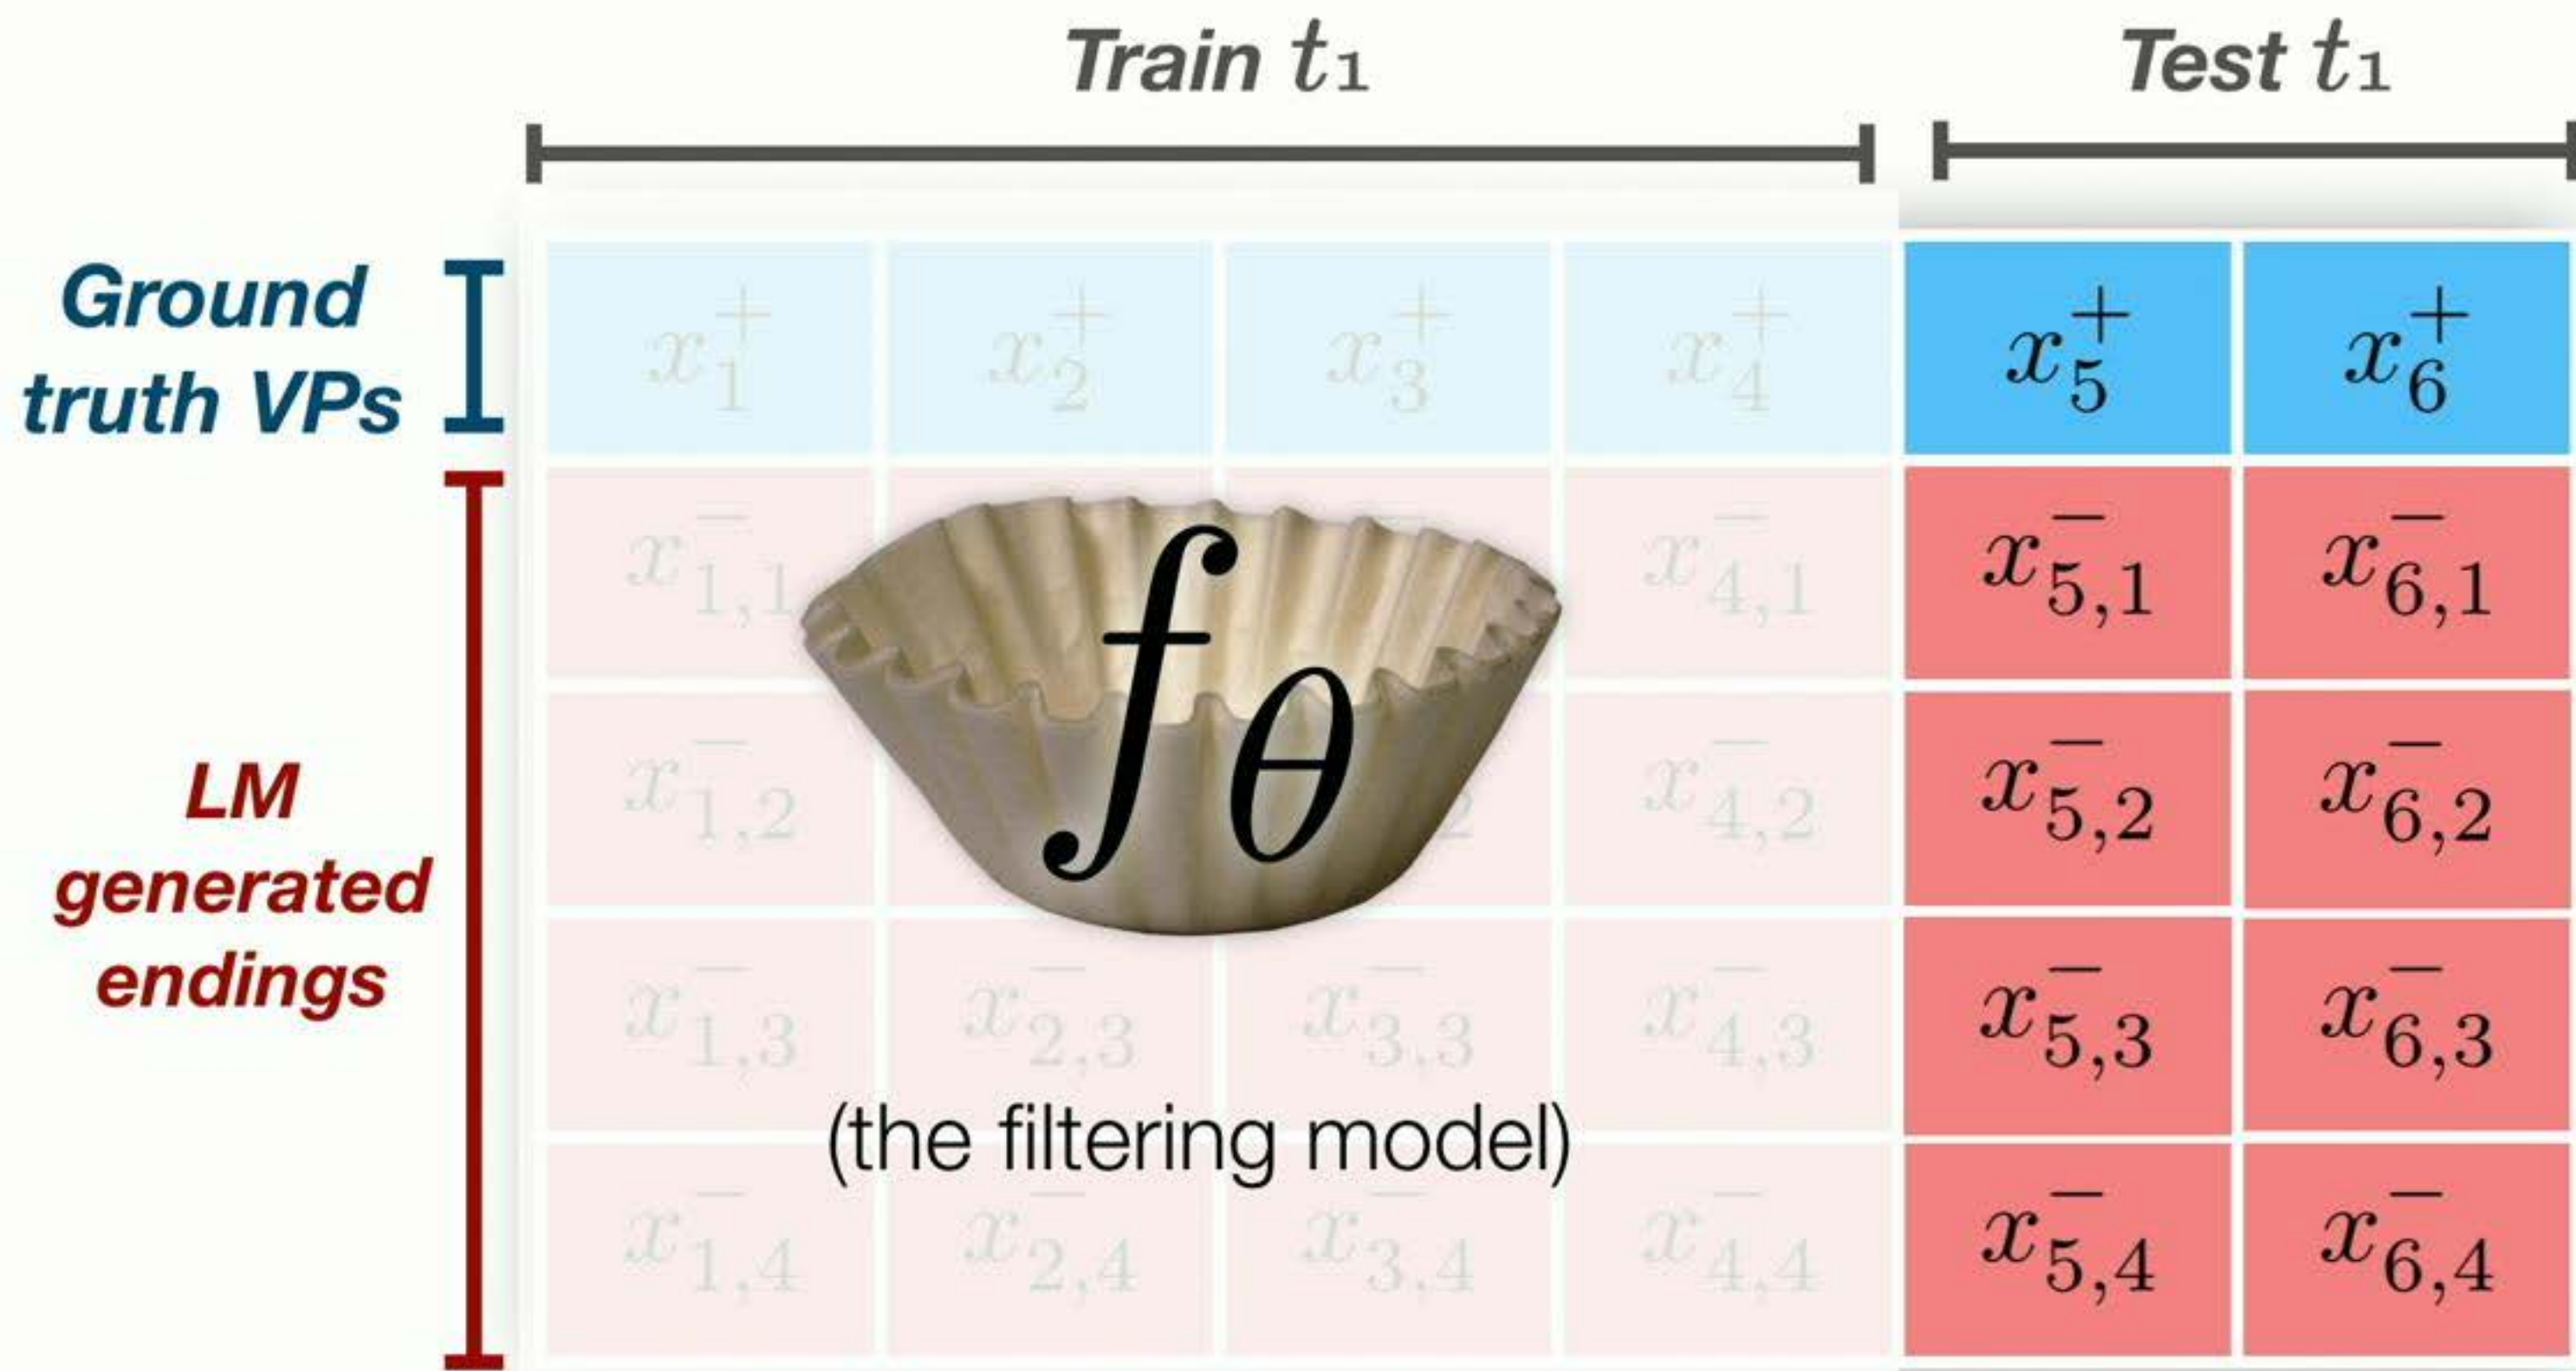
⋮

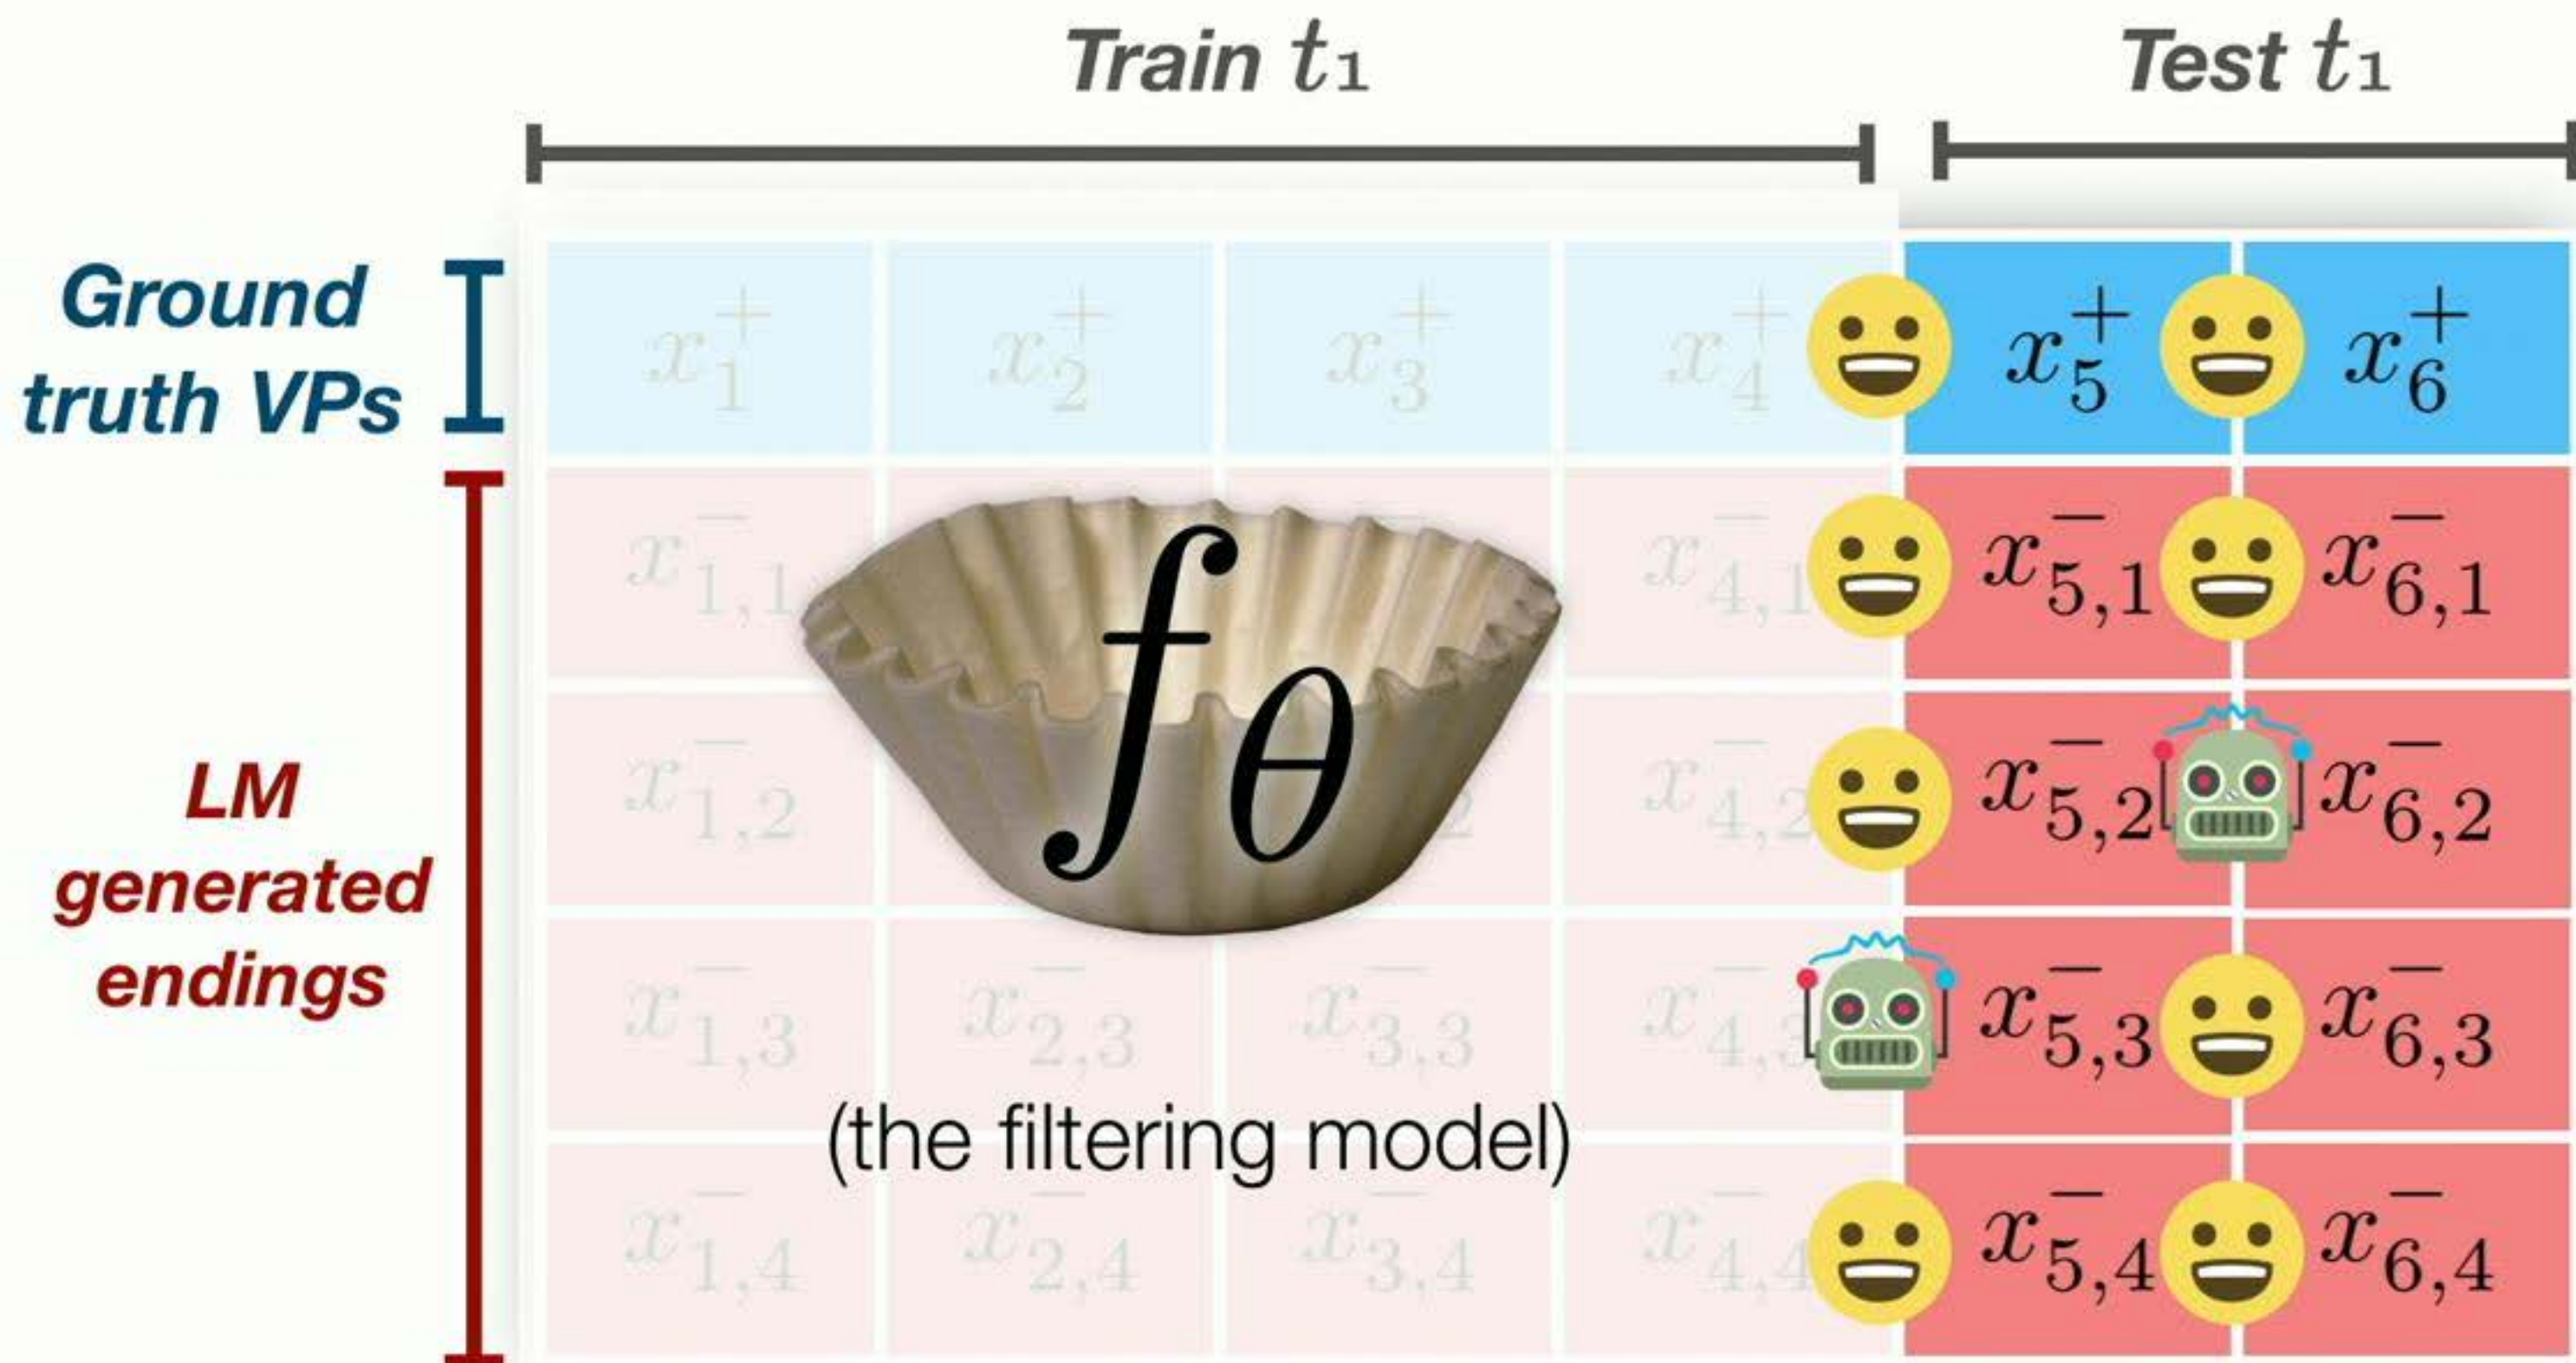
⋮

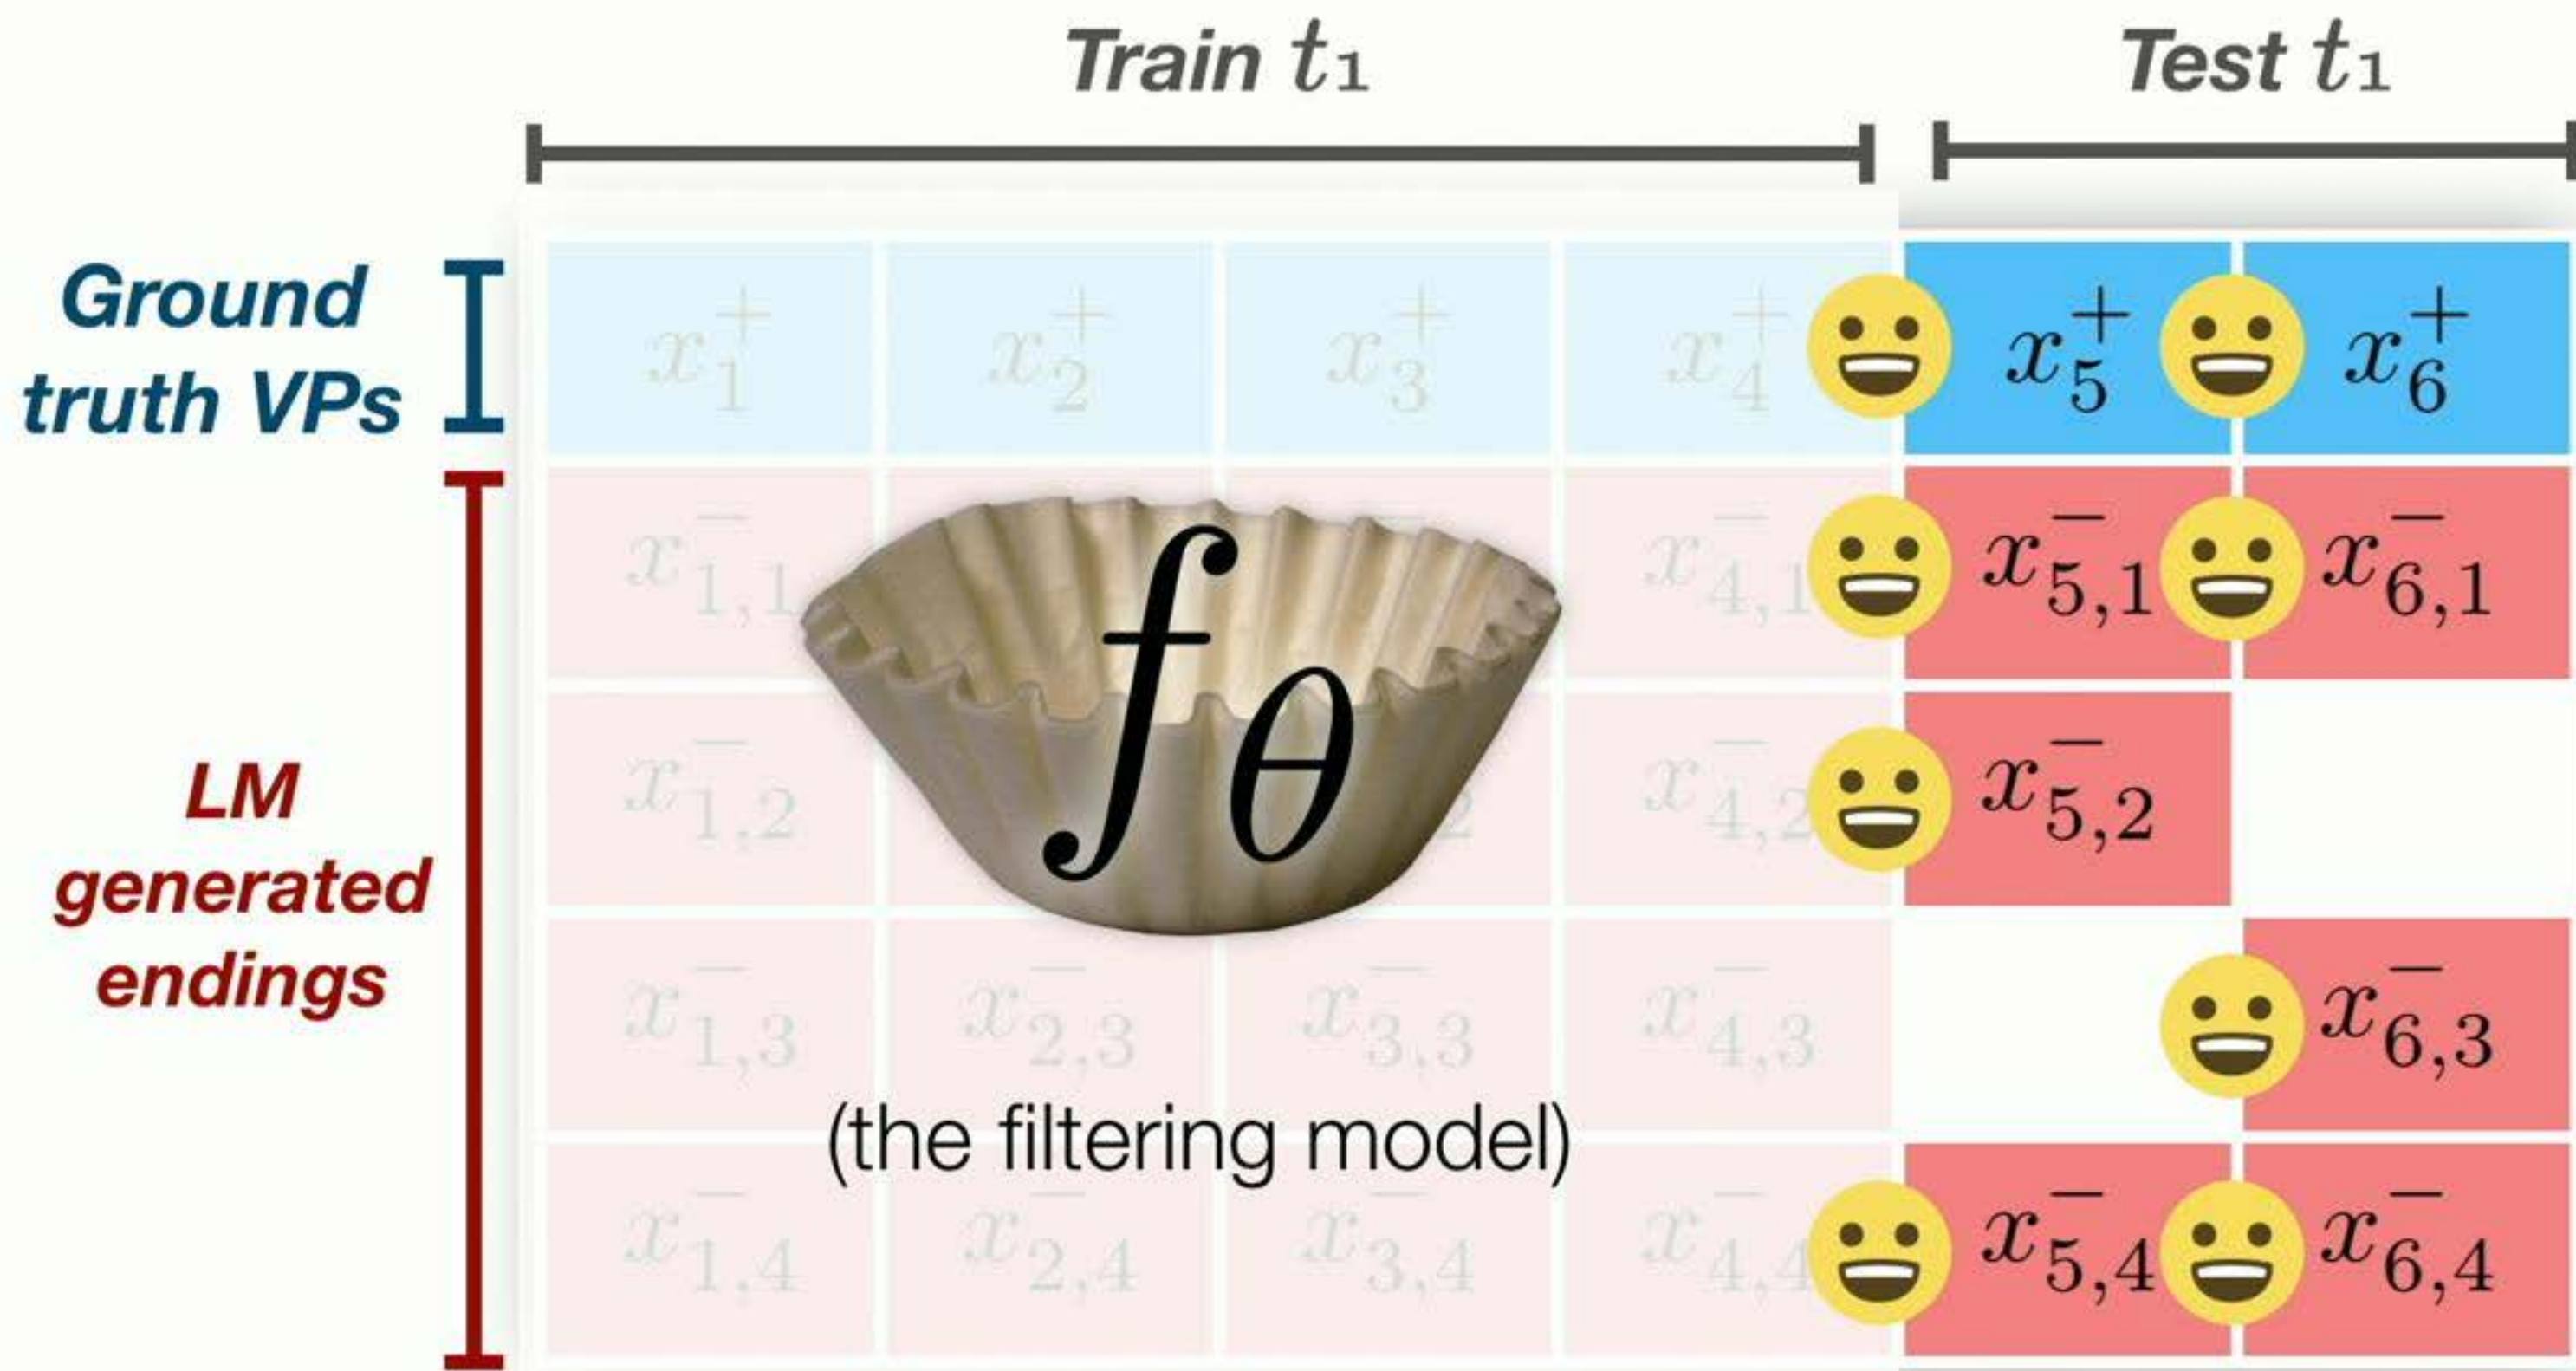
⋮

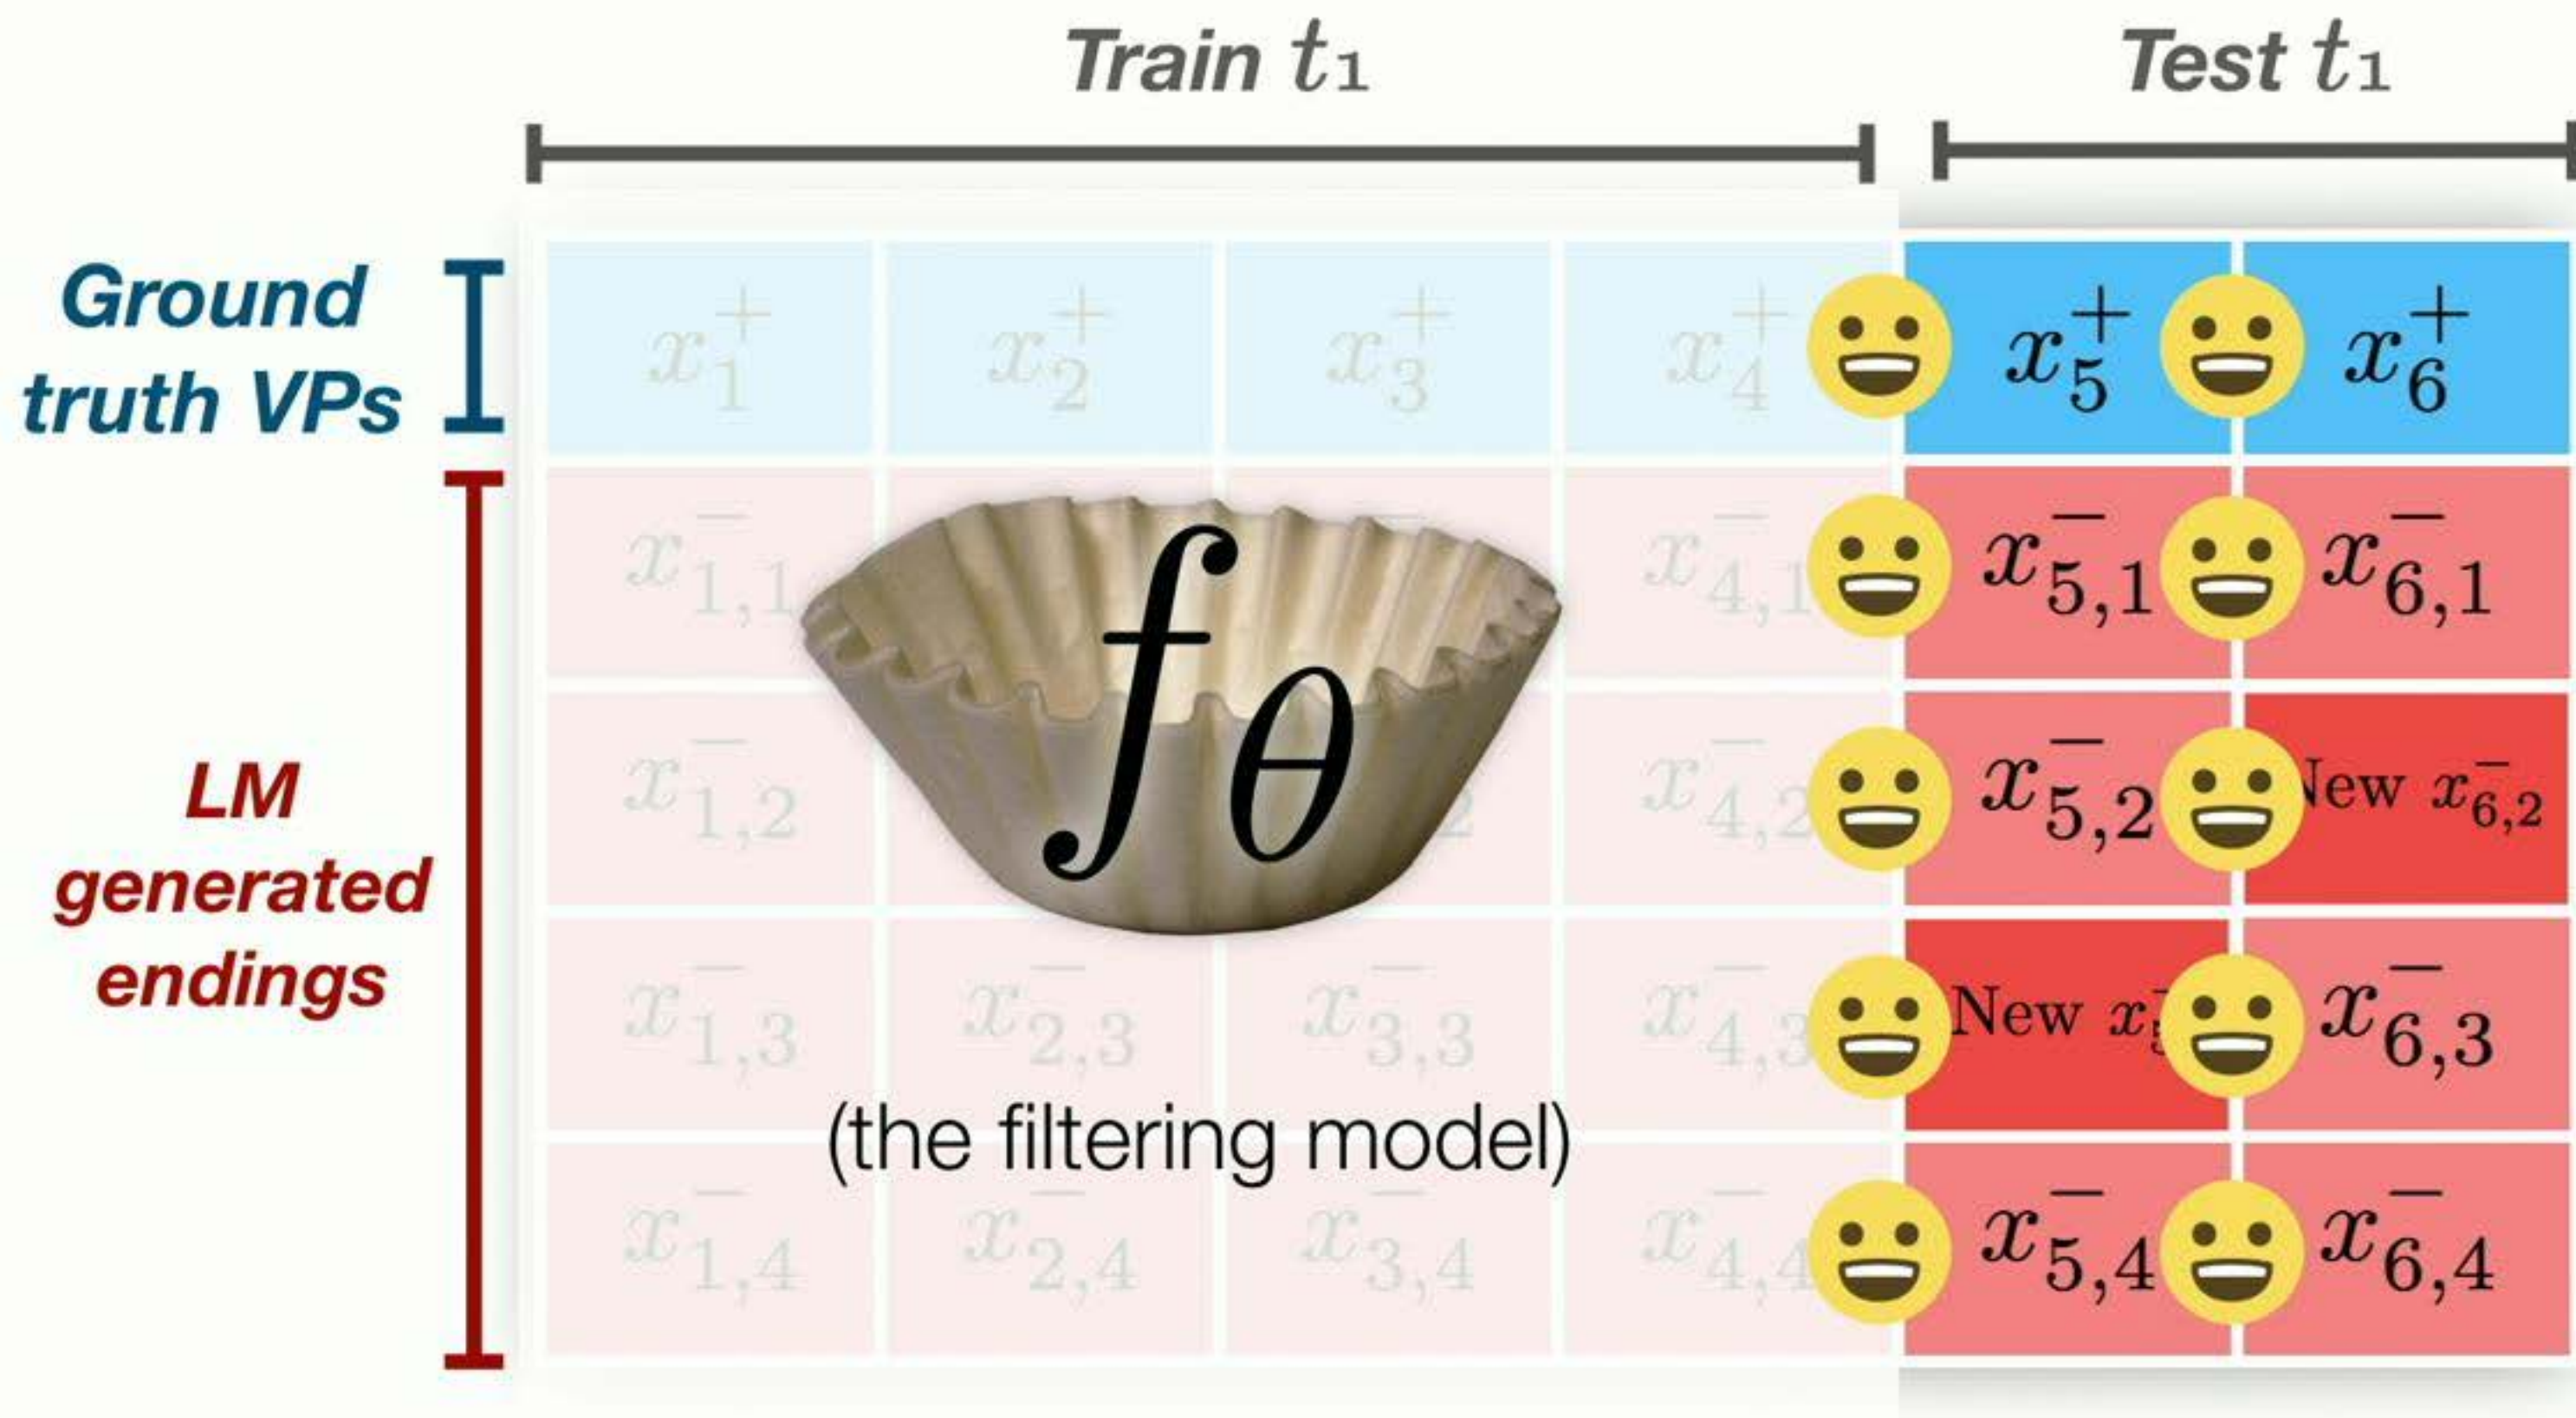
**LM
generated
endings**

<i>Train t_1</i>						<i>Test t_1</i>	
Ground truth VPs	I	x_1^+	x_2^+	x_3^+	x_4^+	x_5^+	x_6^+
		$x_{1,1}^-$	$x_{2,1}^-$	$x_{3,1}^-$	$x_{4,1}^-$	$x_{5,1}^-$	$x_{6,1}^-$
LM generated endings	I	$x_{1,2}^-$	$x_{2,2}^-$	$x_{3,2}^-$	$x_{4,2}^-$	$x_{5,2}^-$	$x_{6,2}^-$
		$x_{1,3}^-$	$x_{2,3}^-$	$x_{3,3}^-$	$x_{4,3}^-$	$x_{5,3}^-$	$x_{6,3}^-$
		$x_{1,4}^-$	$x_{2,4}^-$	$x_{3,4}^-$	$x_{4,4}^-$	$x_{5,4}^-$	$x_{6,4}^-$

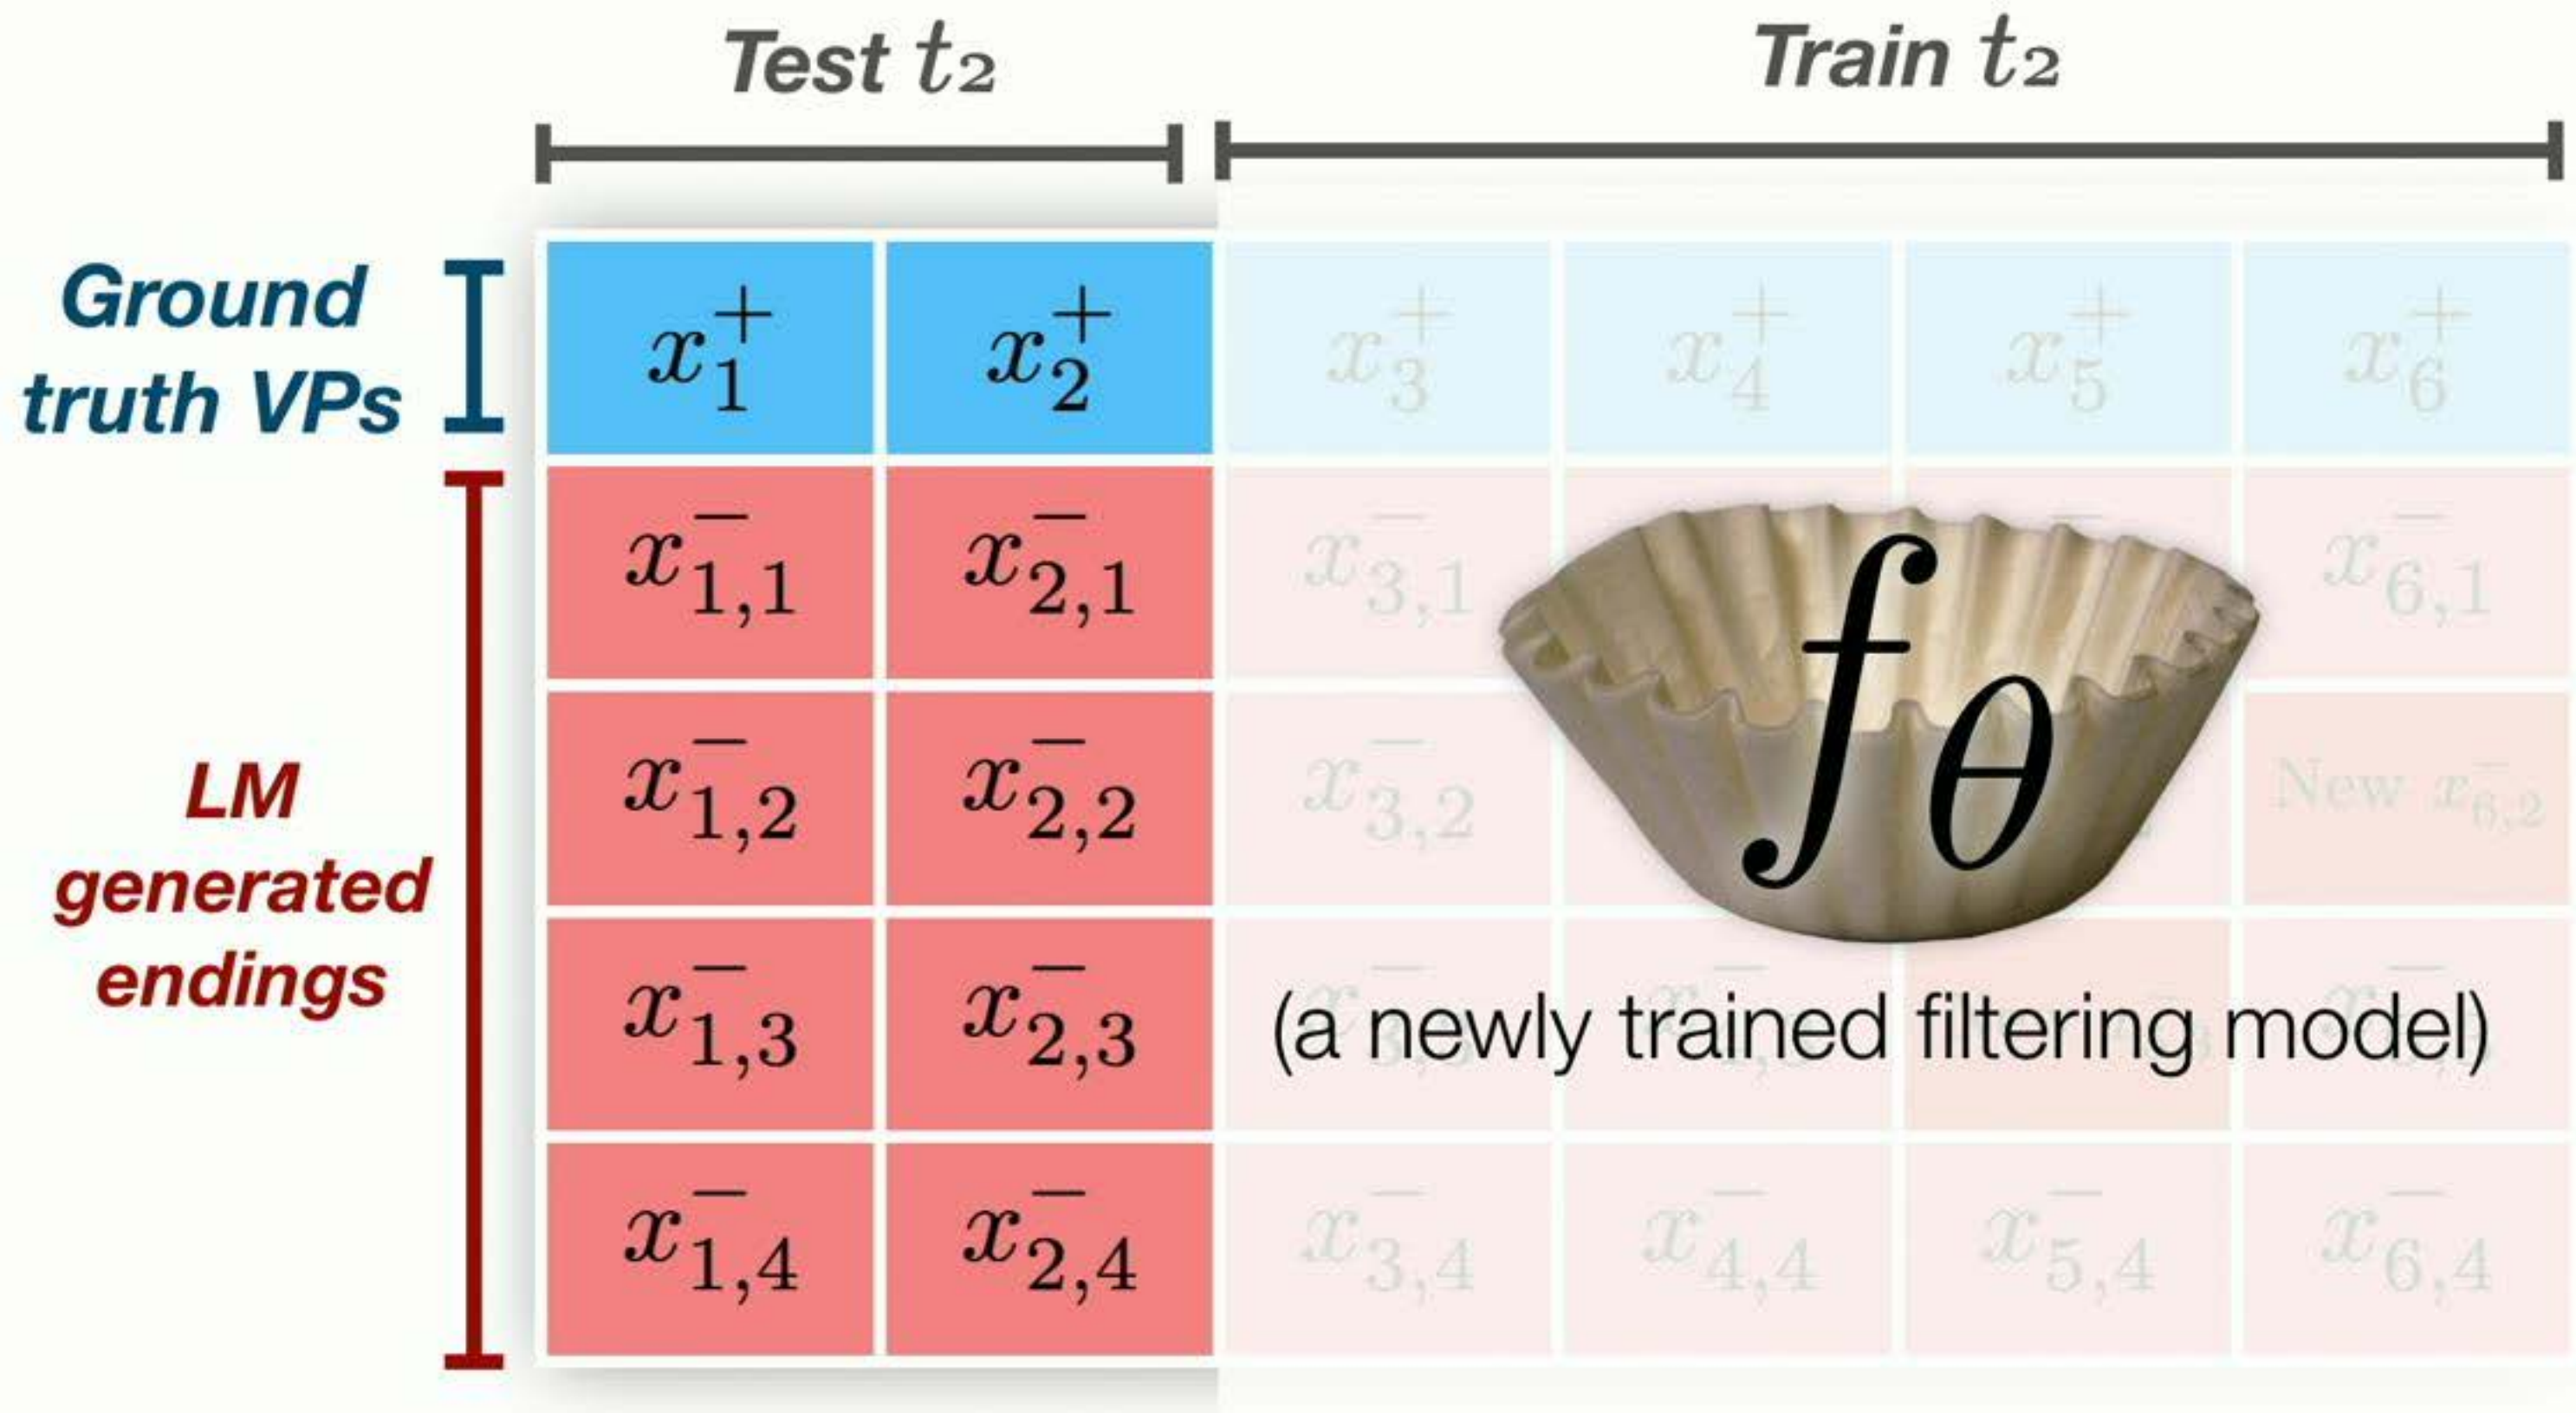


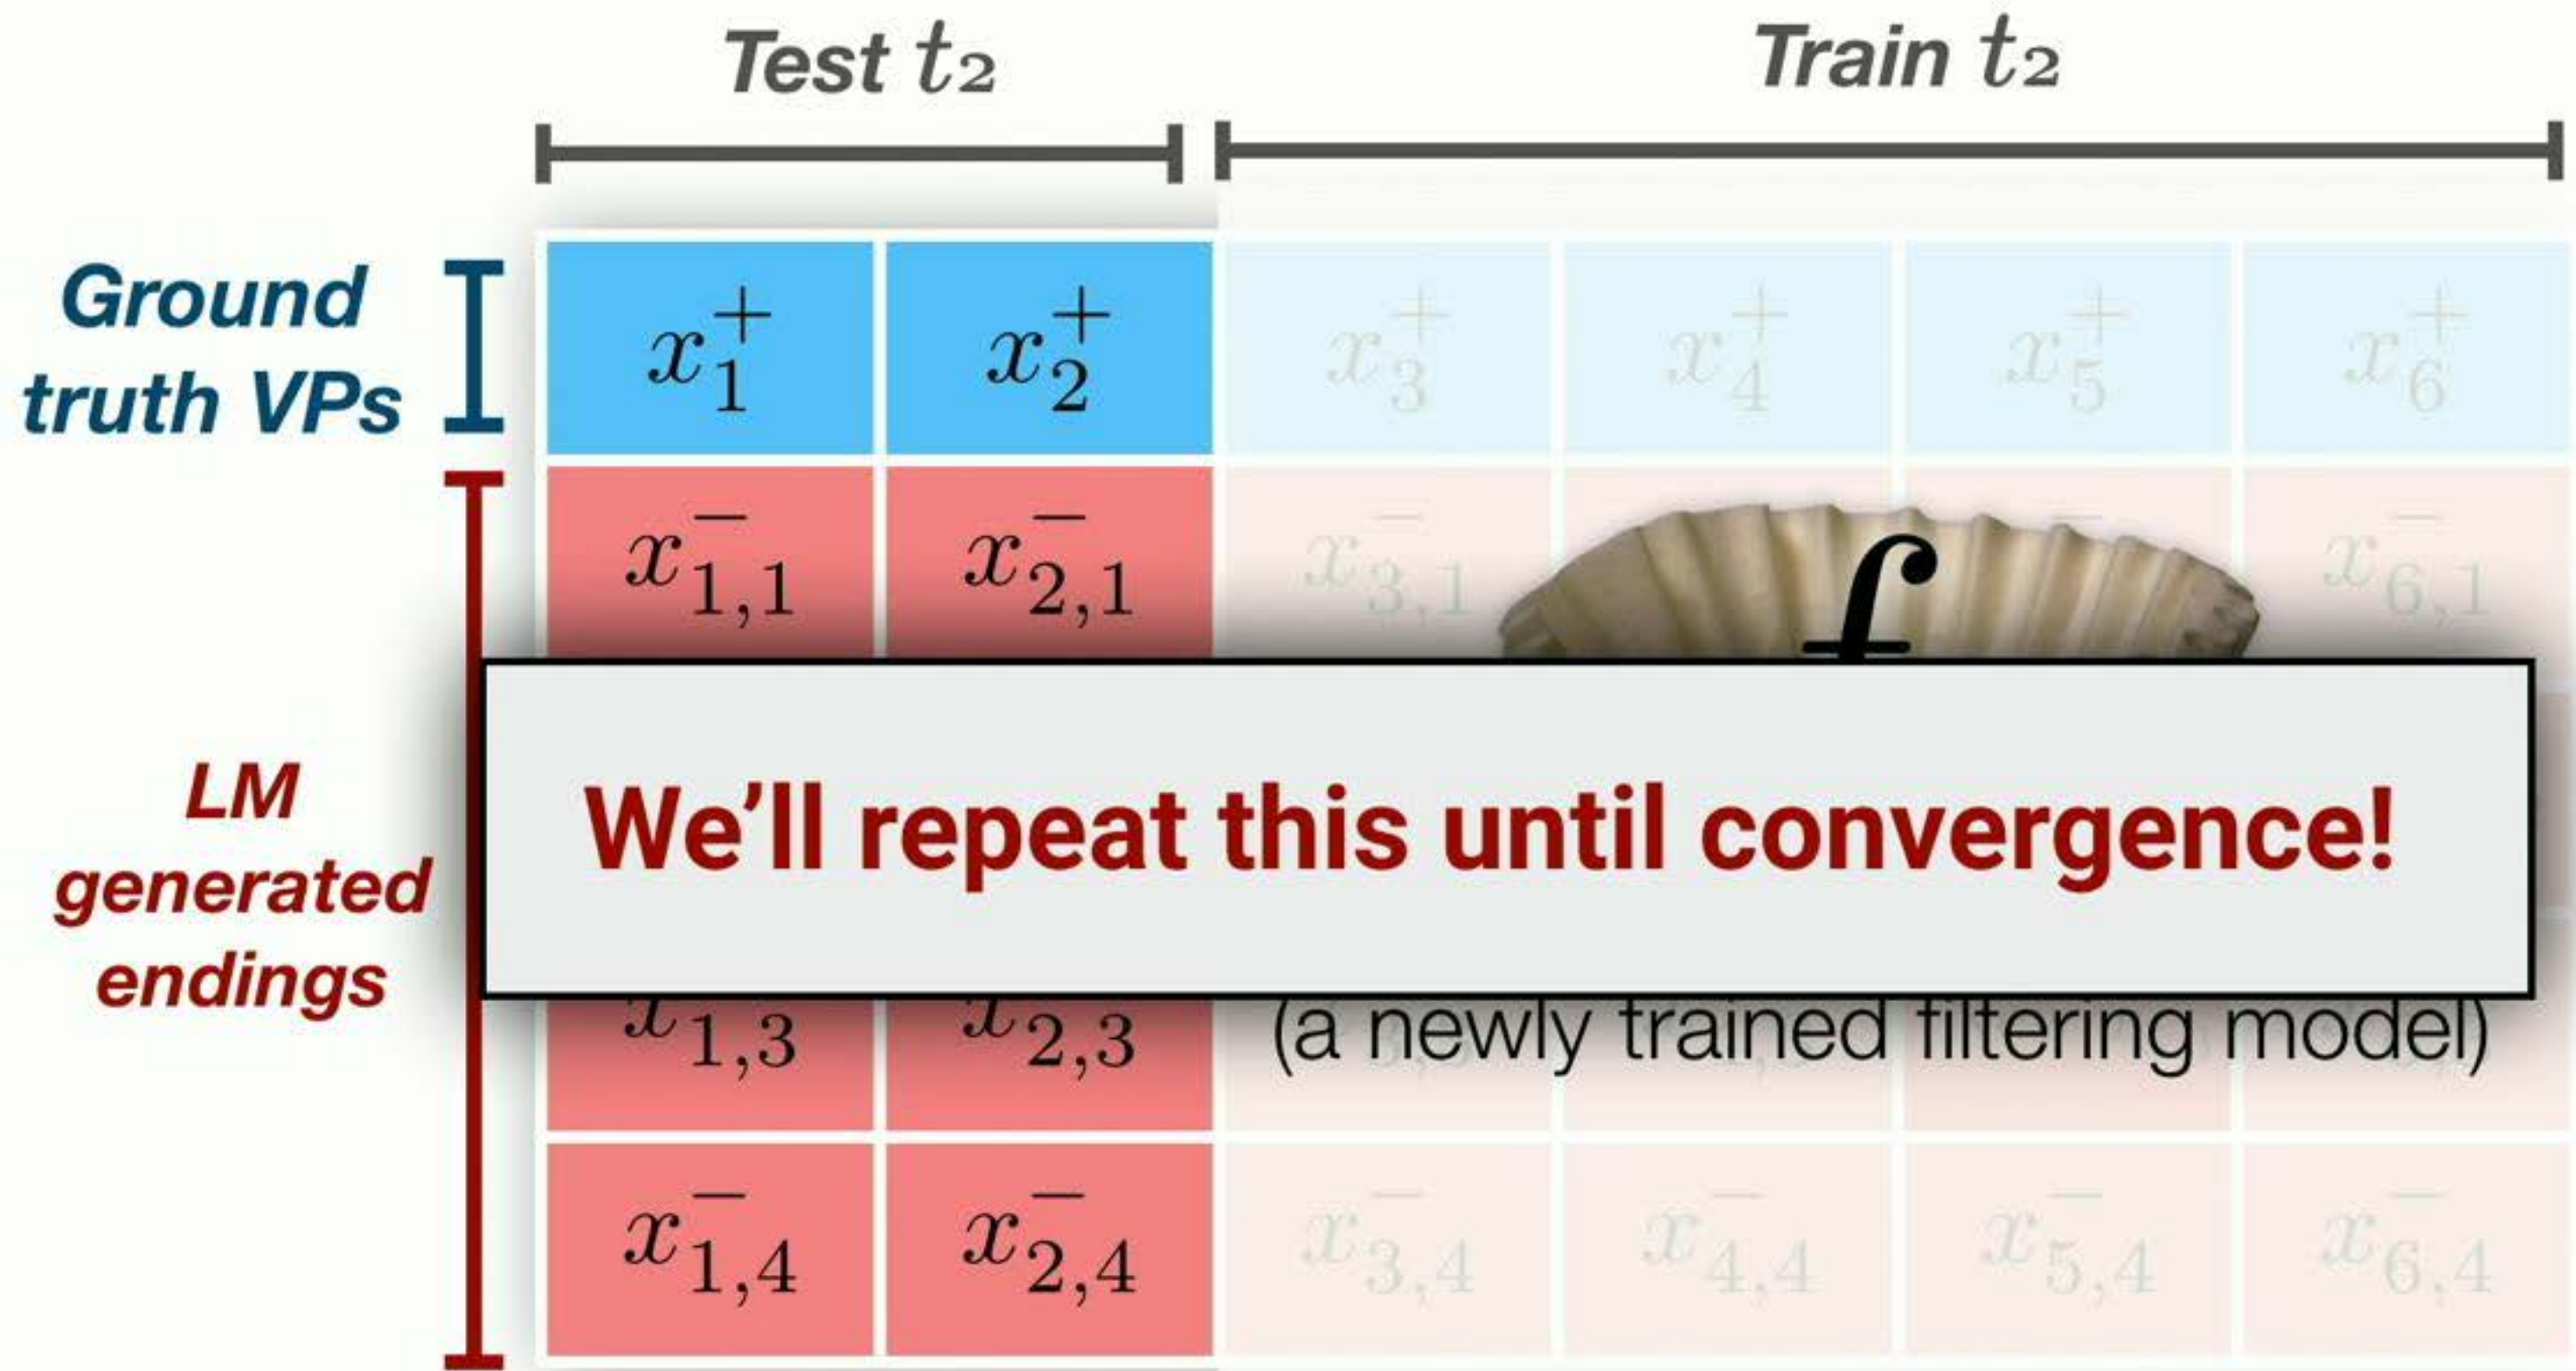







		<i>Test t_2</i>			<i>Train t_2</i>		
Ground truth VPs LM generated endings	I	x_1^+	x_2^+	x_3^+	x_4^+	x_5^+	x_6^+
		$x_{1,1}^-$	$x_{2,1}^-$	$x_{3,1}^-$	$x_{4,1}^-$	$x_{5,1}^-$	$x_{6,1}^-$
		$x_{1,2}^-$	$x_{2,2}^-$	$x_{3,2}^-$	$x_{4,2}^-$	$x_{5,2}^-$	New $x_{6,2}^-$
		$x_{1,3}^-$	$x_{2,3}^-$	$x_{3,3}^-$	$x_{4,3}^-$	New $x_{5,3}^-$	$x_{6,3}^-$
		$x_{1,4}^-$	$x_{2,4}^-$	$x_{3,4}^-$	$x_{4,4}^-$	$x_{5,4}^-$	$x_{6,4}^-$







What AF does and doesn't do

What AF does and doesn't do




What AF does and doesn't do



Easy!

What AF does and doesn't do




What AF does and doesn't do



\neq



?



????????


What AF does and doesn't do



\neq



?



Didn't you
read the disclaimer?



**THIS MODEL SHOULD ONLY
BE EVALUATED ON DATA
THAT'S SIMILAR TO THE
TRAINING SET!!!!**

What AF does and doesn't do



~



?


What AF does and doesn't do



~



?



! ? ! ? !


What AF does and doesn't do



~



?



This is our goal with AF!!!!

Math behind AF

Dataset $\mathcal{D} = \{(x_i, y_i)\}_{1 \leq i \leq N}$ Model $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$

Loss function $L(f_\theta, \mathcal{D})$

Math behind AF

Dataset $\mathcal{D} = \{(x_i, y_i)\}_{1 \leq i \leq N}$ Model $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$

Loss function $L(f_\theta, \mathcal{D})$

Empirical
error

$$I(\mathcal{D}, f) = \frac{1}{N} \sum_{i=1}^N L(f_{\theta_i^*}, \{(x_i, y_i)\}),$$

where $\theta_i^* = \operatorname{argmin}_{\theta} L(f_\theta, \mathcal{D} \setminus \{(x_i, y_i)\})$,

Math behind AF

Dataset $\mathcal{D} = \{(x_i, y_i)\}_{1 \leq i \leq N}$ Model $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$

Loss function $L(f_\theta, \mathcal{D})$

Empirical
error

$$I(\mathcal{D}, f) = \frac{1}{N} \sum_{i=1}^N L(f_{\theta_i^*}, \{(x_i, y_i)\}),$$

where $\theta_i^* = \underset{\theta}{\operatorname{argmin}} L(f_\theta, \mathcal{D} \setminus \{(x_i, y_i)\}),$

The best parameters as determined by other datapoints

Math behind AF

Dataset $\mathcal{D} = \{(x_i, y_i)\}_{1 \leq i \leq N}$ Model $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^{|\mathcal{Y}|}$

How much does knowing N-1 datapoints help?

Empirical
error

$$I(\mathcal{D}, f) = \frac{1}{N} \sum_{i=1}^N L(f_{\theta_i^*}, \{(x_i, y_i)\}),$$

where $\theta_i^* = \underset{\theta}{\operatorname{argmin}} L(f_\theta, \mathcal{D} \setminus \{(x_i, y_i)\}),$

The best parameters as determined by other datapoints

Math behind AF

Let's **MAXIMIZE** dataset-level empirical error!

\mathcal{Y}

How much does knowing N-1 datapoints help?

Empirical
error

$$I(\mathcal{D}, f) = \frac{1}{N} \sum_{i=1}^N L(f_{\theta_i^*}, \{(x_i, y_i)\}),$$

where $\theta_i^* = \underset{\theta}{\operatorname{argmin}} L(f_{\theta}, \mathcal{D} \setminus \{(x_i, y_i)\}),$

The best parameters as determined by other datapoints

Adversarial Filters for SWAG



We used an ensemble of stylistic models centered around the ending sentence. These are powerful adversaries!

(Schwartz et al., 2017, Gururangan et al., 2018)

Adversarial Filters for SWAG



We used an ensemble of stylistic models centered around the ending sentence. These are powerful adversaries!

(Schwartz et al., 2017, Gururangan et al., 2018)



A multilayer perceptron, given LM perplexities as features



Bag-of-words



CNN



BiLSTM, with uncommon words replaced by POS tags

Adversarial Filters for SWAG



We used an ensemble of stylistic models centered around the ending sentence. These are powerful adversaries!

(Schwartz et al., 2017, Gururangan et al., 2018)



After 100 iterations of adversarial filtering, the accuracy of these stylistic models drops to chance.

Adversarial Filters for SWAG



We used an ensemble of stylistic models centered around the ending sentence. These are powerful adversaries!

(Schwartz et al., 2017, Gururangan et al., 2018)




After 100 iterations of adversarial filtering, the accuracy of these stylistic models drops to chance.





Last, crowd workers validate the entire dataset to remove false negatives, leaving us with 110k multiple-choice questions.

Unique Contributions of Adversarial Filtering

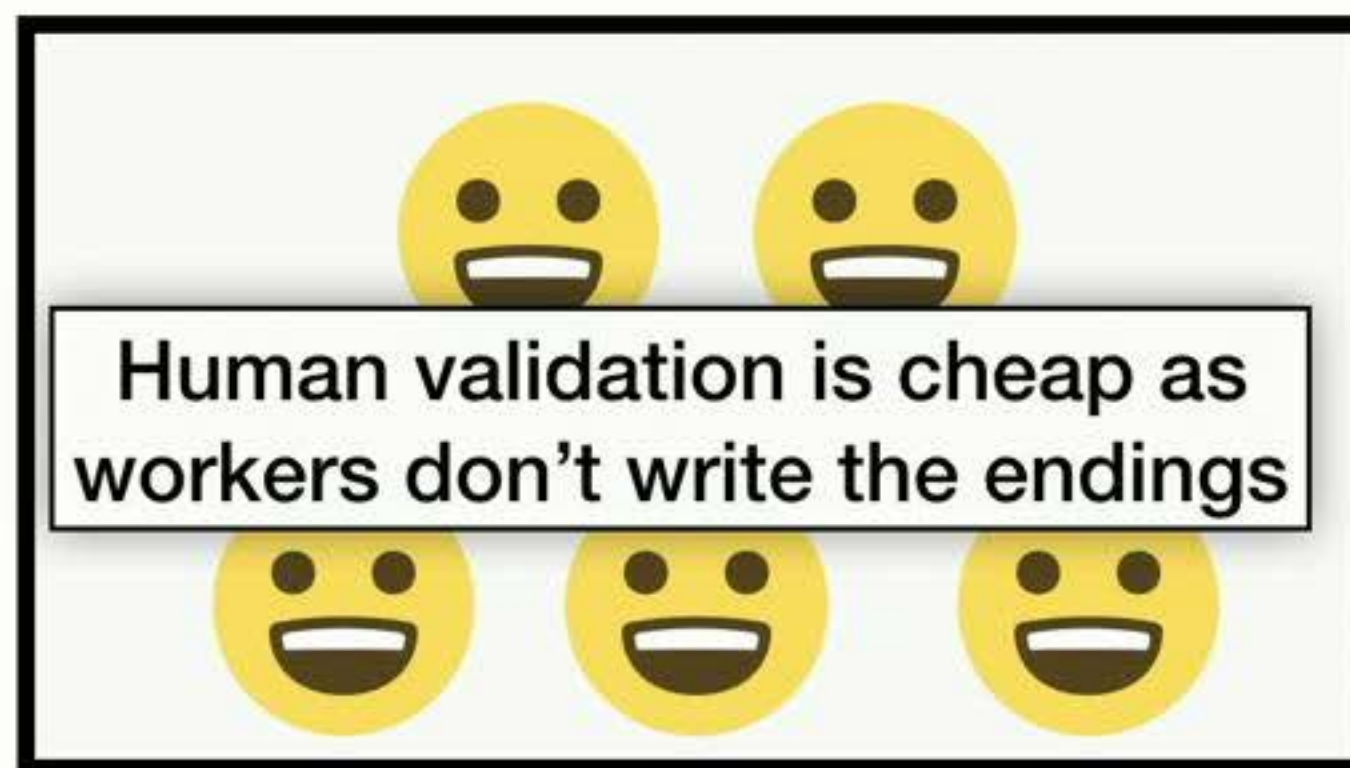


Unique Contributions of Adversarial Filtering




Unique Contributions of Adversarial Filtering




Unique Contributions of Adversarial Filtering



SWAG Results




SWAG Results




SWAG Results




SWAG Results




SWAG Results




SWAG Results




SWAG Results




SWAG Results



What's next?



OpenAI GPT performance vs. training set size



What do models learn by fine-tuning on SWAG?



A girl is going across the monkey bars.



She gets to the other side and stands on a wooden plank.

What do models learn by fine-tuning on SWAG?



A girl is going across the monkey bars.



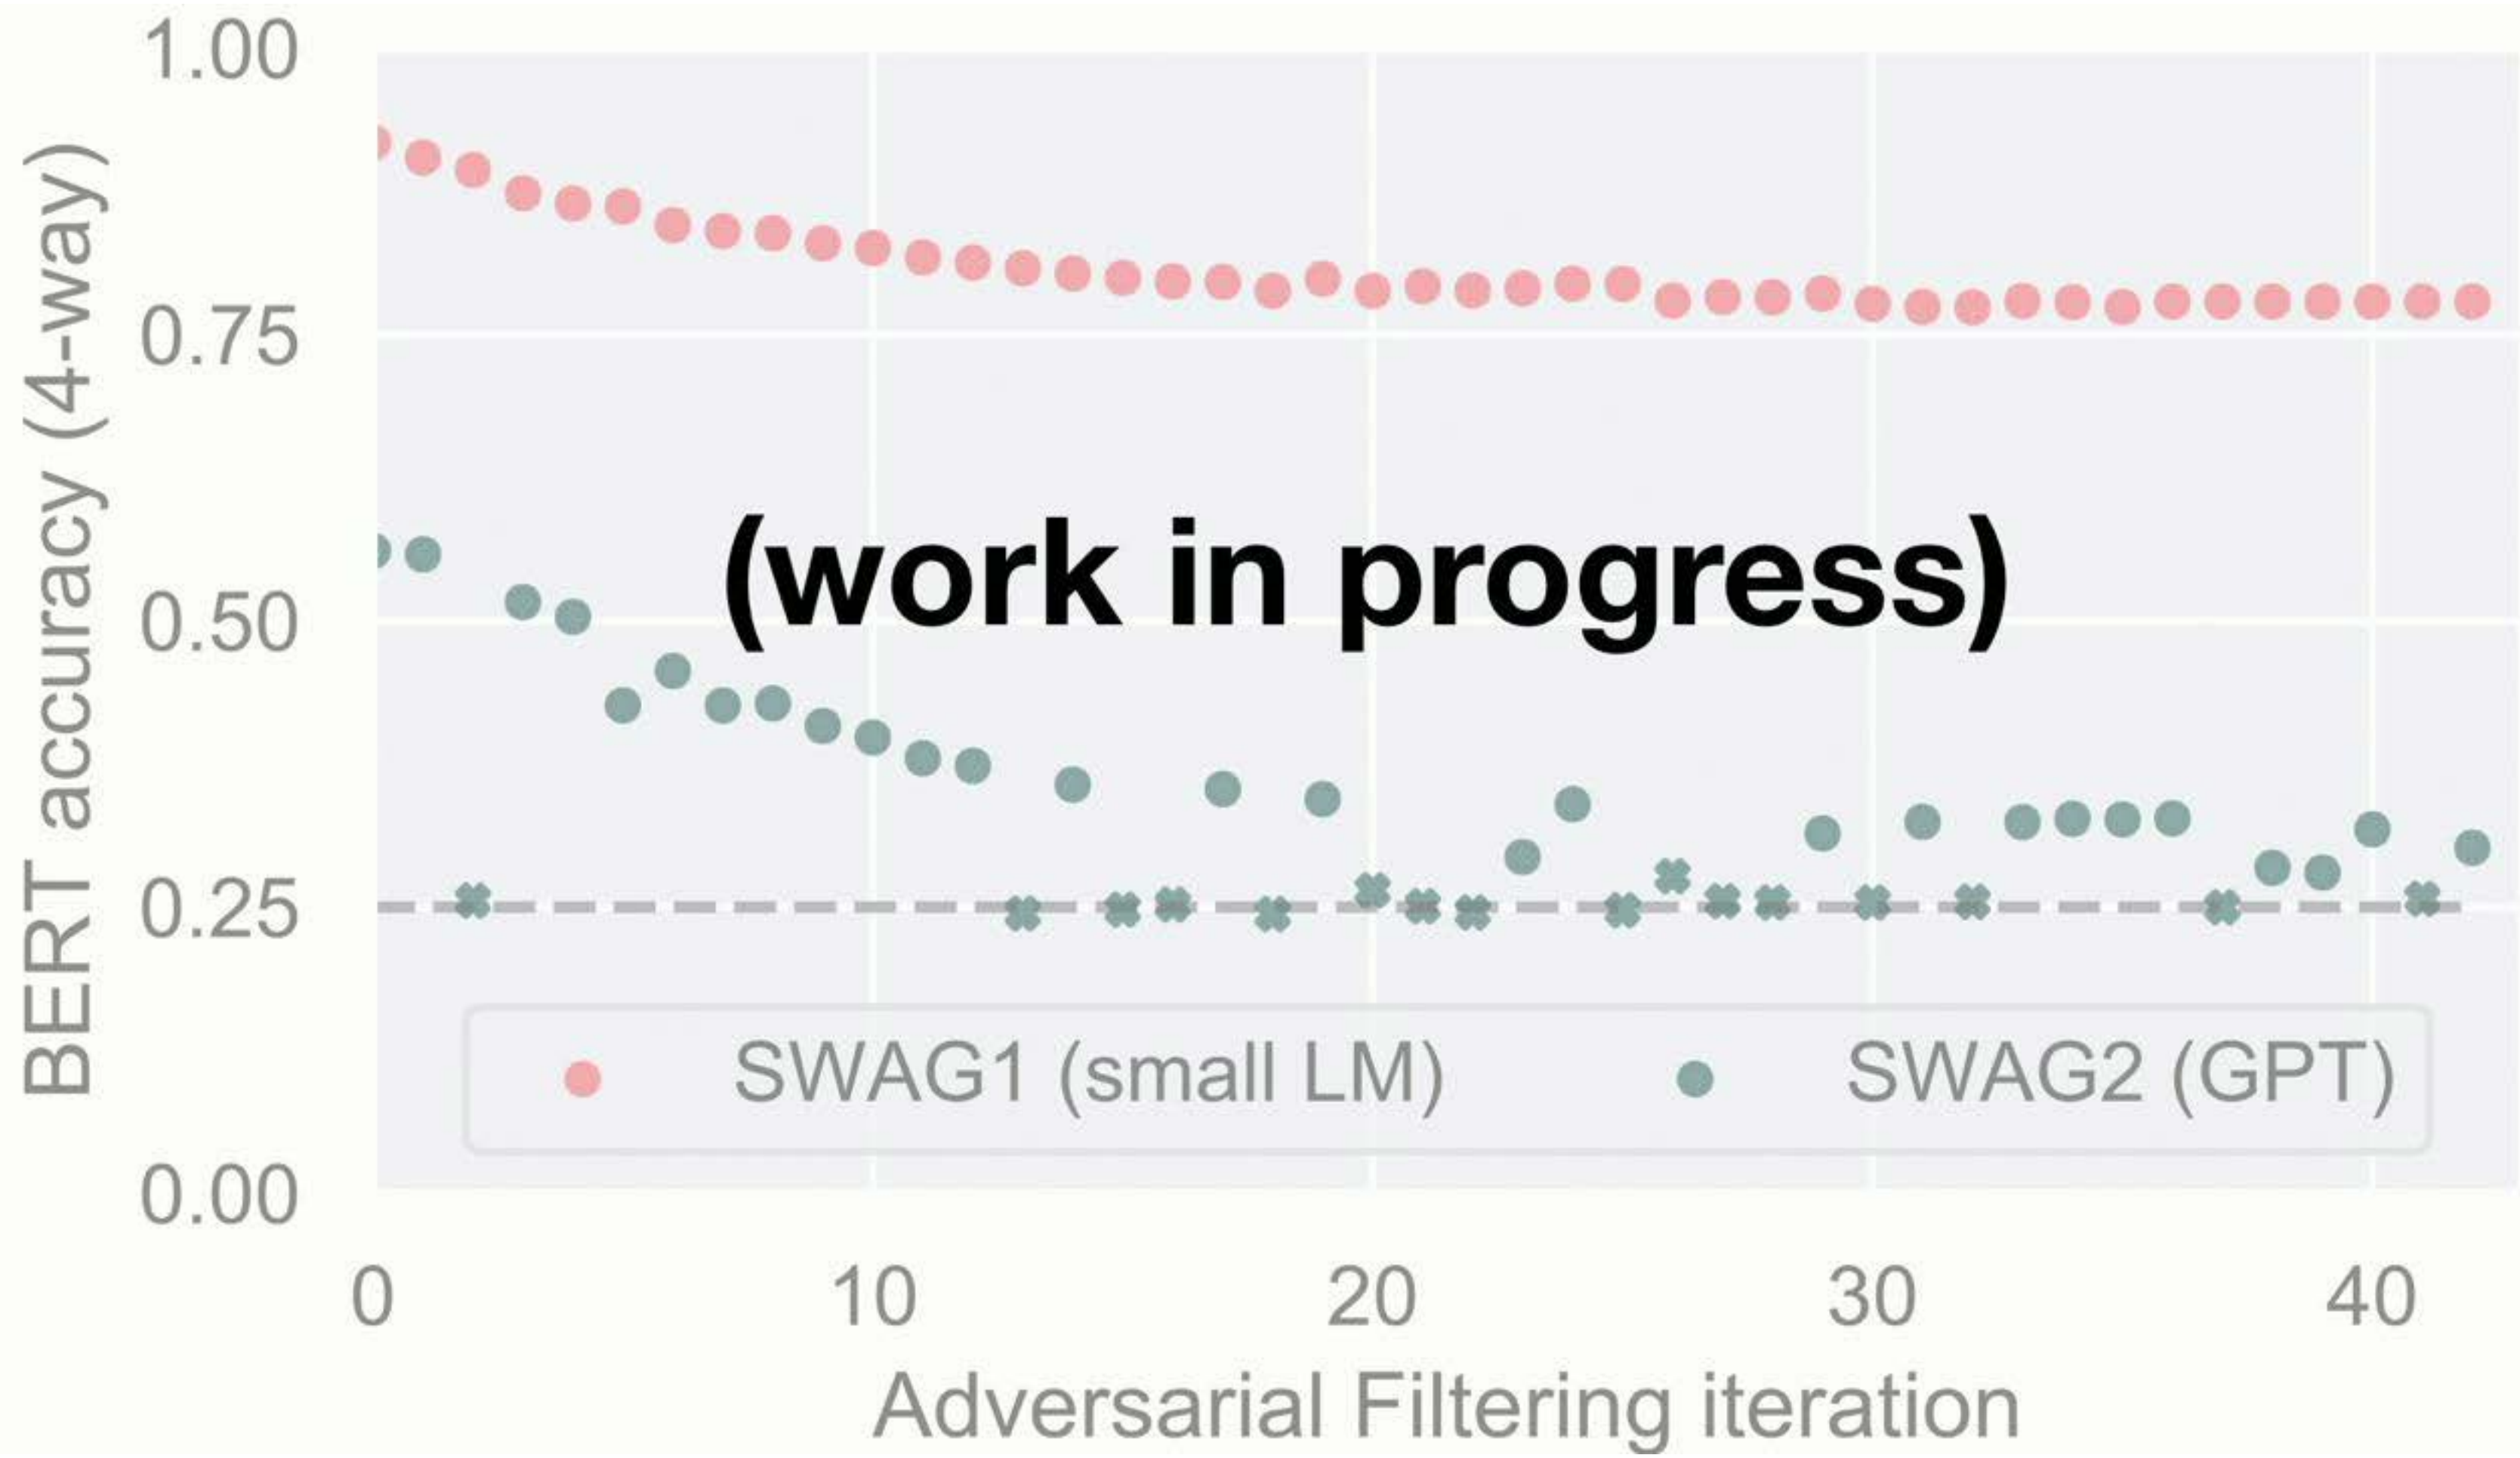
She gets to the other side and stands on a wooden plank.



Can a 2 sentence NLI dataset be constructed at scale while being resistant to pretraining approaches?


Can a 2 sentence NLI dataset be constructed at scale while being resistant to pretraining approaches?


Actually, yes.




Someone is pointing at
someone else. The waiter...







Why is [person4 ] pointing at [person1 ]?




How can we get vision systems to learn this??







Why is [person4 ] pointing at [person1 ]?

From Recognition to Cognition: Visual Commonsense Reasoning (arxiv18)






Why is [person4 ] pointing at [person1 ]?





- a) He is telling [person3 ] that [person1 ] ordered the pancakes.
- b) He just told a joke.
- c) He is feeling accusatory towards [person1 ].
- d) He is giving [person1 ] directions.





Many AI systems
perform well,
but do so for
questionable
reasons





Why is [person4 ] pointing at [person1 ]?

- a) He is telling [person3 ] that [person1 ] ordered the pancakes.
- b) He just told a joke.
- c) He is feeling accusatory towards [person1 ].
- d) He is giving [person1 ] directions.

- a) [person1 ] has the pancakes in front of him.
- b) [person4 ] is taking everyone's order and asked for clarification.
- c) [person3 ] is looking at the pancakes and both she and [person2 ] are smiling slightly.

- d) [person3 ] is delivering food to the table, and she might not know whose order is whose.

*I chose **a)**
because...*




Why is

How can we collect challenging commonsense inferences?

cakes.

a)

b)

c) He is feeling accusatory towards [person1 ].

d)

He is giving [person2 ].

How do we get the wrong answers?

pancakes in

everyone's

g at the pa

are smiling slightly.

How do we model this?

d) [person3 ] is delivering food to the table, and she might not know whose order is whose.



Our contributions

- New task: Visual Commonsense Reasoning





Our contributions

- New task: Visual Commonsense Reasoning



- Building VCR, feat. Adversarial Matching



- Recognition to Cognition Networks





Our contributions

- New task: Visual Commonsense Reasoning



- Building VCR, feat. Adversarial Matching



- Recognition to Cognition Networks



Collecting commonsense inferences

Google commonsense reasoning

All Images News Videos Shopping More Settings Tools

distributional winograd schema trinh event calculus distributional semantics commonst

Commonsense reasoning - Wikipedia
en.wikipedia.org

Common Sense Reasoning for Inter...
ocw.mit.edu



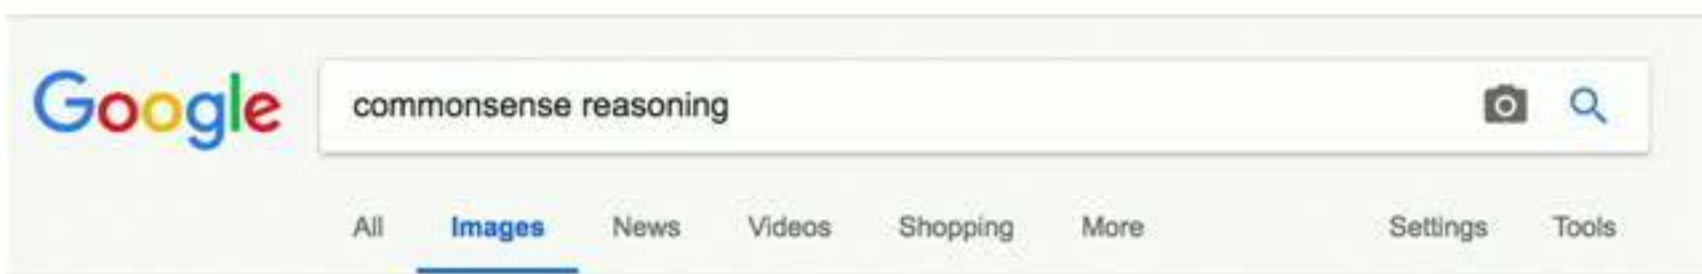
Commonsense Reasoning and Commonsense ...
m-cacm.acm.org

Common Sense Reasoning ...
web.media.mit.edu

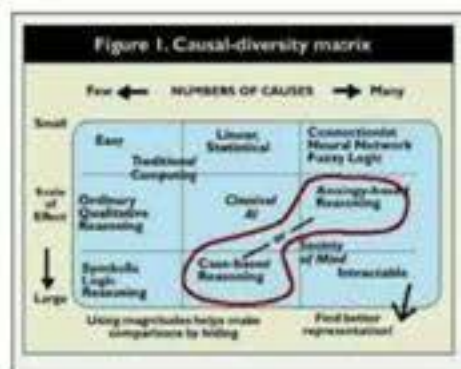
Common Sense Reasoning Projects - Fall 2006
web.media.mit.edu

combining commonsense reasoning ...
researchgate.net


Collecting commonsense inferences



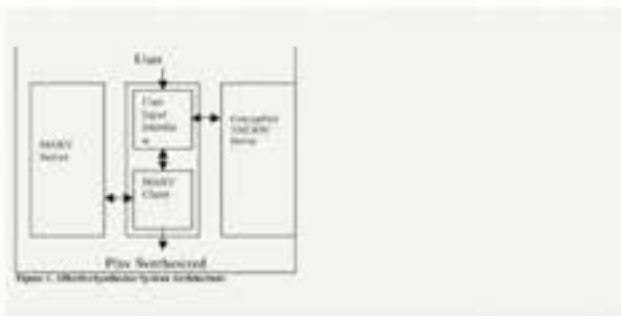
Commonsense reasoning - Wikipedia
en.wikipedia.org




Common Sense Reasoning for Inter...
ocw.mit.edu

Commonsense Reasoning and Commonsense ...
m-cacm.acm.org


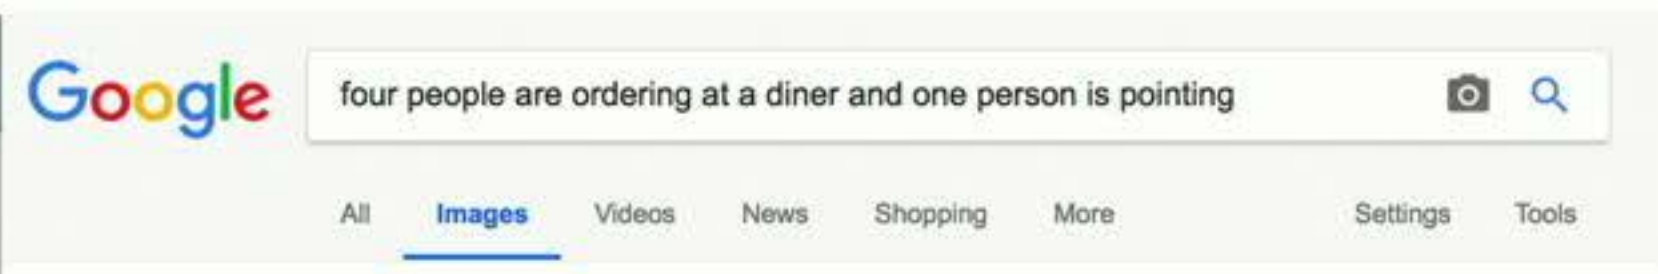
Common Sense Reasoning ...
web.media.mit.edu



Common Sense Reasoning Projects - Fall 2006
web.media.mit.edu



combining commonsense reasoning ...
researchgate.net



Eating Out: A Basic Guide to Restaur...
tokyocheapo.com



Going off-menu in a restaurant can ...
independent.co.uk



Impact Your Restaurant Order ...
glamour.com



The Postmates Problem: Why Some ...
eater.com




Order Food at a Japanese Restaurant ...
kawakawalearningstudio.com



How Restaurants Can Deal...
eater.com

Collecting commonsense inferences



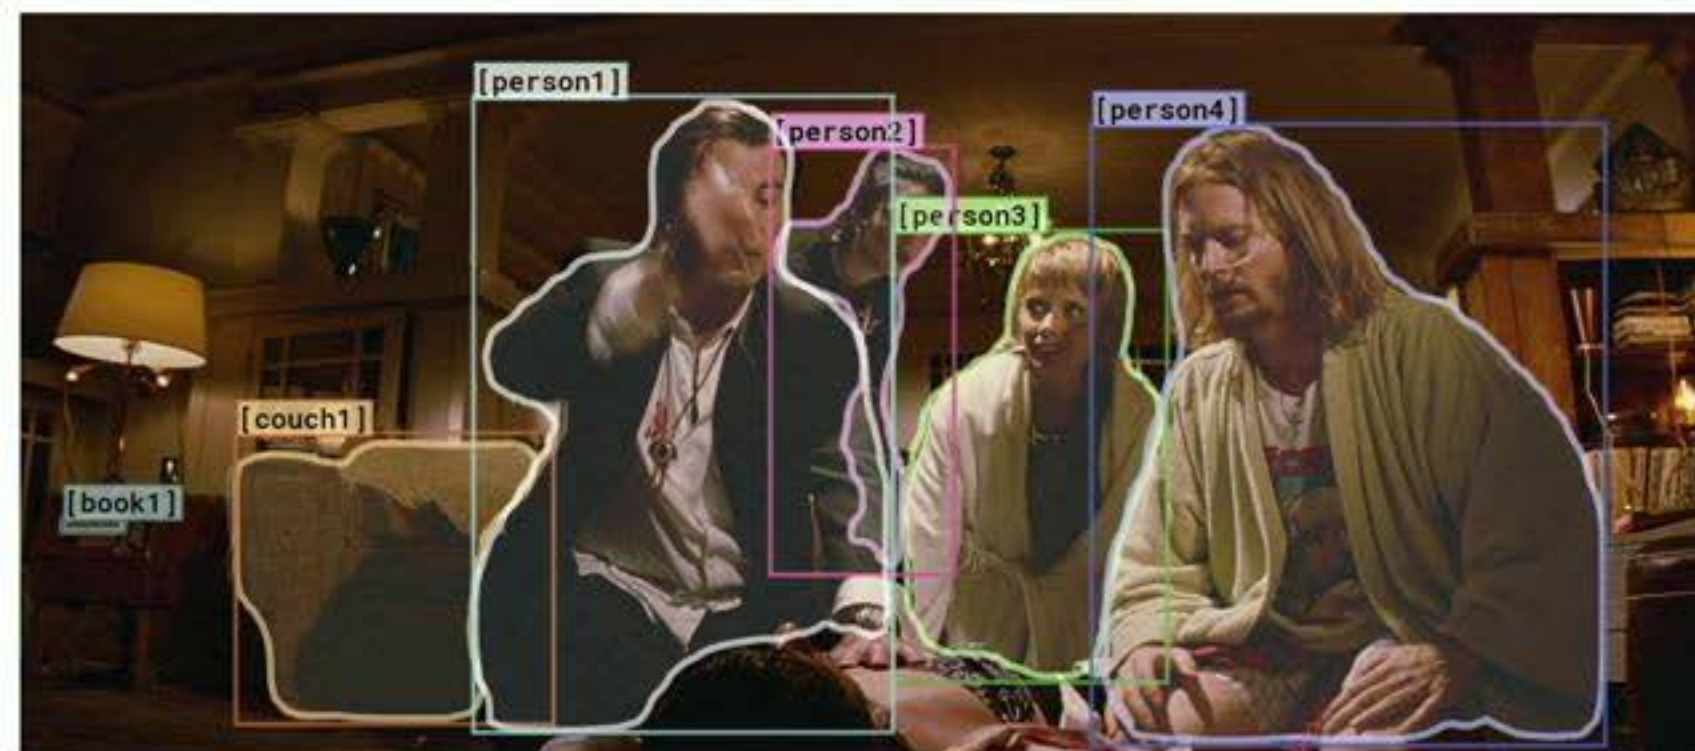
Collecting commonsense inferences



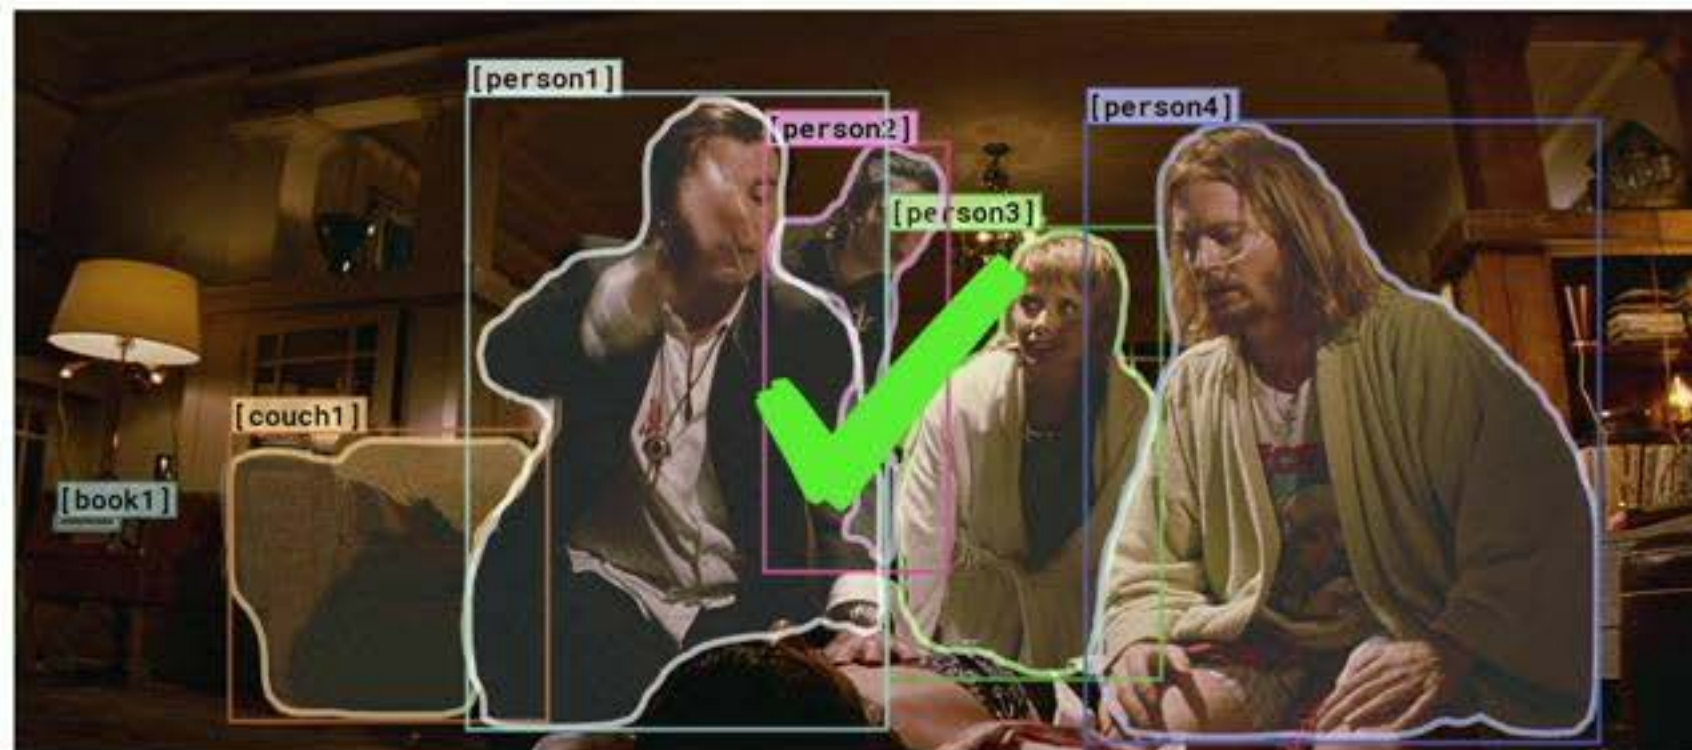
Collecting commonsense inferences



Collecting commonsense inferences



Collecting commonsense inferences



Collecting commonsense inferences

Examples (see rowanzellers.com/halloffame for more) [\(expand/collapse\)](#)

Past caption: The driver looks out and sees the Ducati's front wheel waggling about in midair.

This caption: The bike's front wheel drops down, and SOMEONE speeds through town.

Next caption: SOMEONE weaves precariously through traffic and hops onto a pavement.



hide all

show all

1 (person)

2 (person)

3 (car)

4 (car)

Question
Q1

Answer
A1

Rationale
R1

- Unlikely (<25% chance)
- Possible (25% to 75%)
- Likely (>75% chance)


Collecting commonsense inferences

Examples (see rowanzellers.com/halloffame for more) [\(expand/collapse\)](#)

Past caption: The driver looks out and sees the Ducati's front wheel wagging about in midair.

This caption: The bike's front wheel drops down, and SOMEONE speeds through town.

Next caption: SOMEONE weaves precariously through traffic and hops onto a pavement.



hide all

show all

1 (person)

2 (person)

3 (car)

4 (car)

Q1

Wh

A1

Answer


R1

Rationale

- Unlikely (<25% chance)
- Possible (25% to 75%)
- Likely (>75% chance)

Collecting commonsense inferences

Next caption: SOMEONE weaves precariously through traffic and hops onto a pavement.



hide all show all 1 (person) 2 (person) 3 (car) 4 (car)

Q1	Where was 1 previously?	A1	1 was in the hospital.	R1	1 is wearing a hospital gown.	<input type="radio"/> Unlikely (<25% chance) <input type="radio"/> Possible (25% to 75%) <input checked="" type="radio"/> Likely (>75% chance)
Q2	Question (optional)	A2	Answer (optional)	R2	Rationale (optional)	<input type="radio"/> Unlikely (<25% chance) <input type="radio"/> Possible (25% to 75%) <input checked="" type="radio"/> Likely (>75% chance)

**How do we get the wrong answers,
avoiding annotation artifacts?**

what color is the shirt?

what color is the shirt?

teal

blue

violet

magenta

Adversarial Matching

Wrong answers must have be
relevant to the question yet different from the correct answer.

Adversarial Matching

Wrong answers must have be

relevant to the question yet different from the correct answer.

Q, A' ***Question
Relevance***

A, A' ***Entailment***

Adversarial Matching

Wrong answers must have be

relevant to the question yet different from the correct answer.

Q, A' *Question
Relevance*

A, A' *Entailment*



We'll use these two metrics to **recycle right answers** to other questions, using a minimum weight bipartite matching.




Adversarial Matching

We'll use these two metrics to **recycle right answers** to other questions, using a minimum weight bipartite matching.

Why are [person1] and [person3] holding their foreheads together?

Why do [person1] and [person3] have their hands clasped?



They are about to kiss.

[person1] and [person3] are praying.




Adversarial Matching

We'll use these two metrics to **recycle right answers** to other questions, using a minimum weight bipartite matching.

Why are [person1] and [person3] holding their foreheads together?

Why do [person1] and [person3] have their hands clasped?



They are about to kiss.

[person1] and [person3] are praying.



Adversarial Matching



Wrong answers must have be

relevant to the question

yet

different from the correct answer.

Q, A' ***Question
Relevance***

A, A' ***Entailment***



Adversarial Matching



Wrong answers must have be

relevant to the question

yet

different from the correct answer.

Q, A' **Question
Relevance**

A, A' **Entailment**

Question relevance
dominates - hard
for machines



Entailment
penalty dominates
- easy for humans



Adversarial Matching



Wrong answers must have be

relevant to the question yet different from the correct answer.

Q, A' **Question
Relevance**

A, A' **Entailment**

This works for the rationales too!

Question relevance
dominates - hard
for machines

— Low λ

High λ —

Entailment
penalty dominates
- easy for humans

What about the people tags (`[person5]`)?



Adversarial Matching

Wrong answers must have be

relevant to the question

yet

different from the correct answer.

Q, A'

**Question
Relevance**

A, A'

Entailment



Adversarial Matching



Wrong answers must have be

relevant to the question

yet

different from the correct answer.

Q, A' **Question
Relevance**

A, A' **Entailment**

Question relevance
dominates - hard
for machines



Entailment
penalty dominates
- easy for humans



Adversarial Matching



Wrong answers must have be

relevant to the question yet different from the correct answer.

Q, A' **Question
Relevance**

A, A' **Entailment**

This works for the rationales too!

Question relevance
dominates - hard
for machines

— Low λ

High λ —


Entailment
penalty dominates
- easy for humans

What about the people tags ([person5])?

We'll randomly modify the detection tags in candidate answer to better match the new question/image.

What about the people tags ([person5])?



We'll randomly modify the detection tags in candidate answer to better match the new question/image.



What about the people tags ([person5])?

We'll randomly modify the detection tags in candidate answer to better match the new question/image.

Answer source




He is giving [person3] directions.




What about the people tags ([person5])?

We'll randomly modify the detection tags in candidate answer to better match the new question/image.



He is giving [person3]  directions.



He is giving [person1]  directions.

Unique Contributions of Adversarial Matching

No answer-only bias

0.25

Unique Contributions of Adversarial Matching

No answer-only bias

0.25



Can raise the bar of difficulty during dataset construction

Unique Contributions of Adversarial Matching

No answer-only bias

0.25



Can raise the bar of difficulty during dataset construction

Entailment NLI means that human validation isn't (as) needed



Putting it all together: VCR

- VCR features 290k questions over 110k images, each with answers and rationales.
- The questions are diverse and challenging.






Our contributions

- New task: Visual Commonsense Reasoning
- Building VCR, feat. Adversarial Matching
- Recognition to Cognition Networks





Detour: Setting up the task





Detour: Setting up the task



Question answering

Answer justification



Detour: Setting up the task



Question answering

$$Q \rightarrow A$$

Answer justification

$$QA \rightarrow R$$



Detour: Setting up the task



Question answering

$$\underbrace{Q}_{\text{Query}} \rightarrow A$$

Query

Answer justification

$$\underbrace{QA}_{\text{Query}} \rightarrow R$$

Query


Let's make both tasks have the same format!




Detour: Setting up the task



Question answering



Answer justification




Let's make both tasks have the same format!




Detour: Setting up the task



Question answering



Answer justification




$$Q \rightarrow AR$$


Recognition to Cognition Networks





Recognition to Cognition Networks



Recognition to Cognition Networks



Query


Why is [person4]  pointing at [person1] ?

Objects





...



Recognition to Cognition Networks



Query

Why is [person4]  pointing at [person1] ?


Response Choice

He is telling [person3]  that [person1]  ordered pancakes.


Objects





Recognition to Cognition Networks



Query

Why is [person4]  pointing at [person1] ?

Response Choice


He is telling [person3]  that [person1]  ordered pancakes.

Objects



...

1. Grounding
2. Contextualization
3. Reasoning





Part 1: Grounding



Objects




Query

Why is [person4]  pointing at [person1] ?

Response Choice

He is telling [person3]  that [person1]  ordered pancakes.



Part 1: Grounding


Objects p1 p2 p2 p3 ...

Query

Why is [person4] pointing at [person1]?

Response Choice

He is telling [person3] that [person1] ordered pancakes.



ResNet

Part 1: Grounding

Objects




Query

Why is [person4]
pointing at [person1]?

Response Choice

He is telling [person3]
that [person1]
ordered pancakes.

Part 1: Grounding



ResNet

Objects



Query

Why is [person4] pointing at [person1]?

Response Choice

He is telling [person3] that [person1] ordered pancakes.





Objects





...

Part 1: Grounding

Query

Why is [person4 ] pointing at [person1 ]?

Response Choice

He is telling [person3 ] that [person1 ] ordered pancakes.









Part 1: Grounding



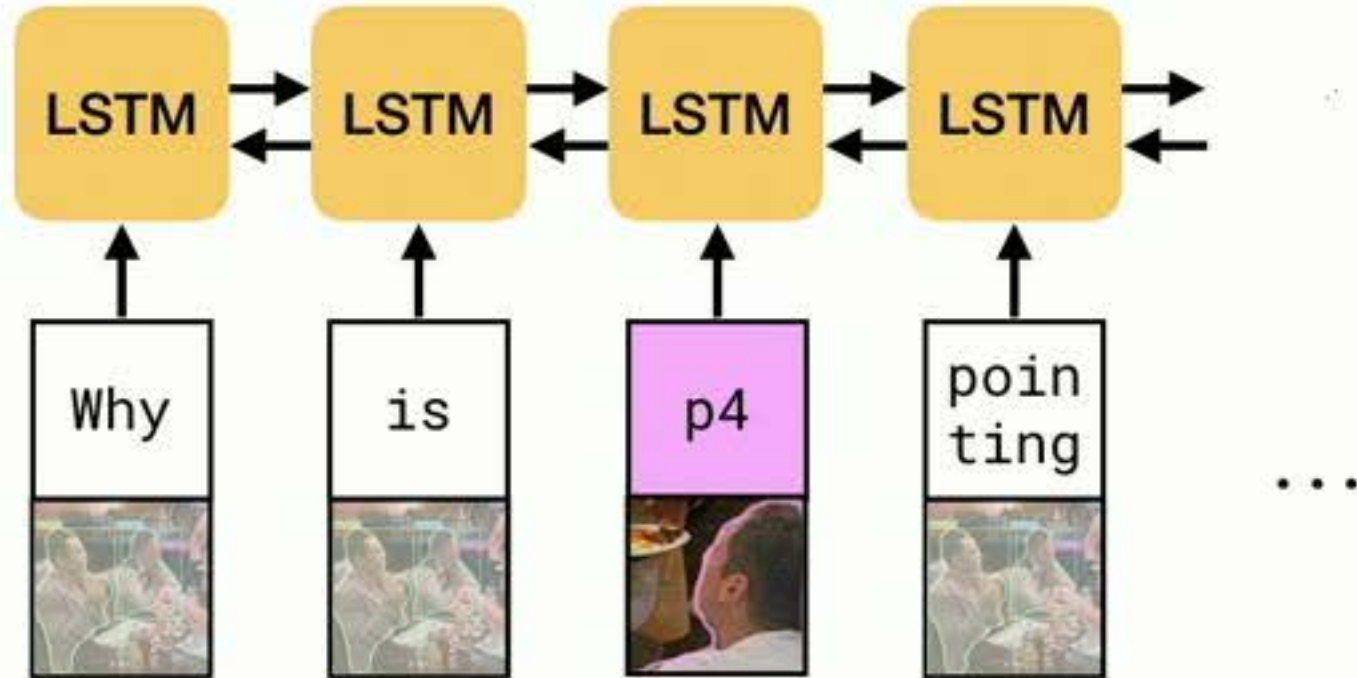
Query

Why is [person4 ] pointing at [person1 ]?



Response Choice

He is telling [person3 ] that [person1 ] ordered pancakes.

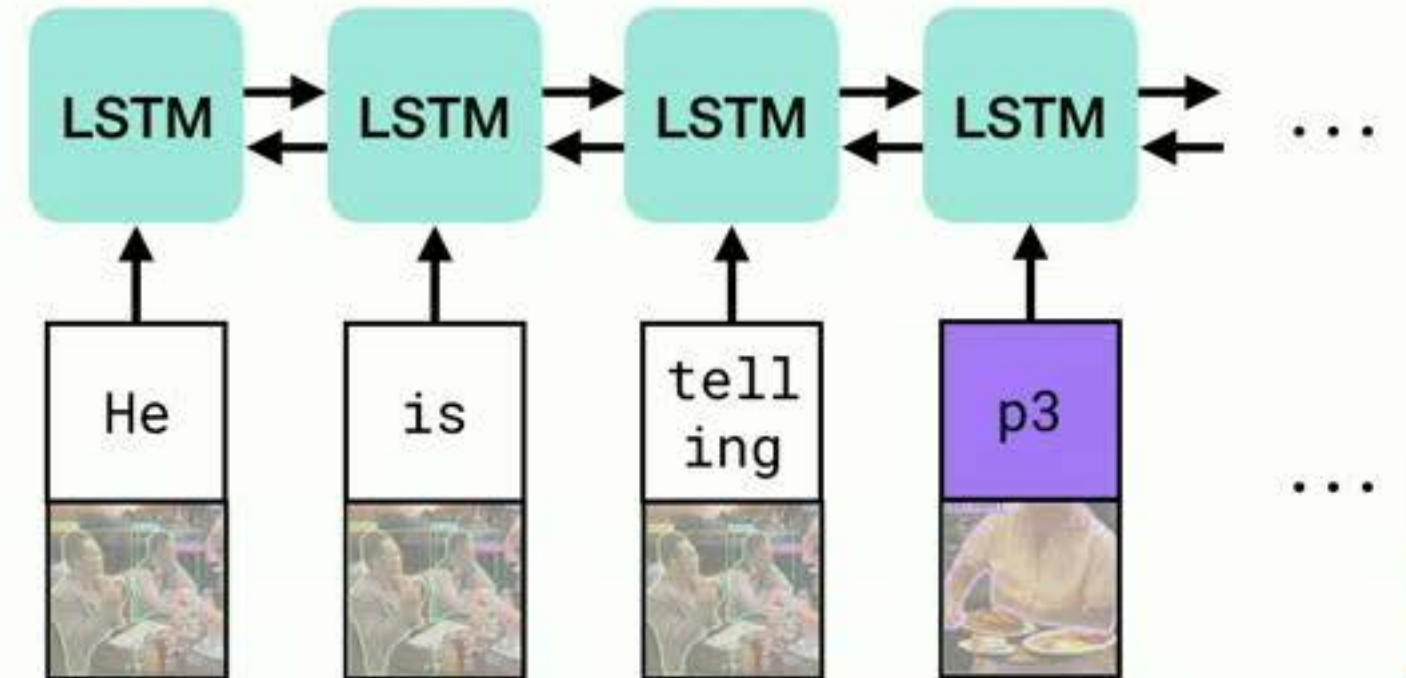






Query

Why is [person4 ] pointing at [person1 ]?

Part 1: Grounding



Response Choice

He is telling [person3 ] that [person1 ] ordered pancakes.







Part 2: Contextualization



Objects



Query

Why is [person4 ] pointing at [person1 ]?

Response Choice


He is telling [person3 ] that [person1 ] ordered pancakes.





Part 2: Contextualization

Objects





	Why	is		...
He				
is				
telling				
...				

Query

Why is [person4 ] pointing at [person1 ]?

Response Choice

He is telling [person3 ] that [person1 ] ordered pancakes.



Part 2: Contextualization

Objects



	Why	is		...
He				
is				
telling				
...				

He				
is				
telling				
...				

Query

Why is **[person4]**
pointing at **[person1]**?

Response Choice

is telling **[person3]**
that **[person1]**
ordered pancakes.



Part 3: Reasoning

Response Choice



is telling **[person3]**



that **[person1]**



ordered pancakes.






Part 3: Reasoning




Response

Attended Query

Attended Objects

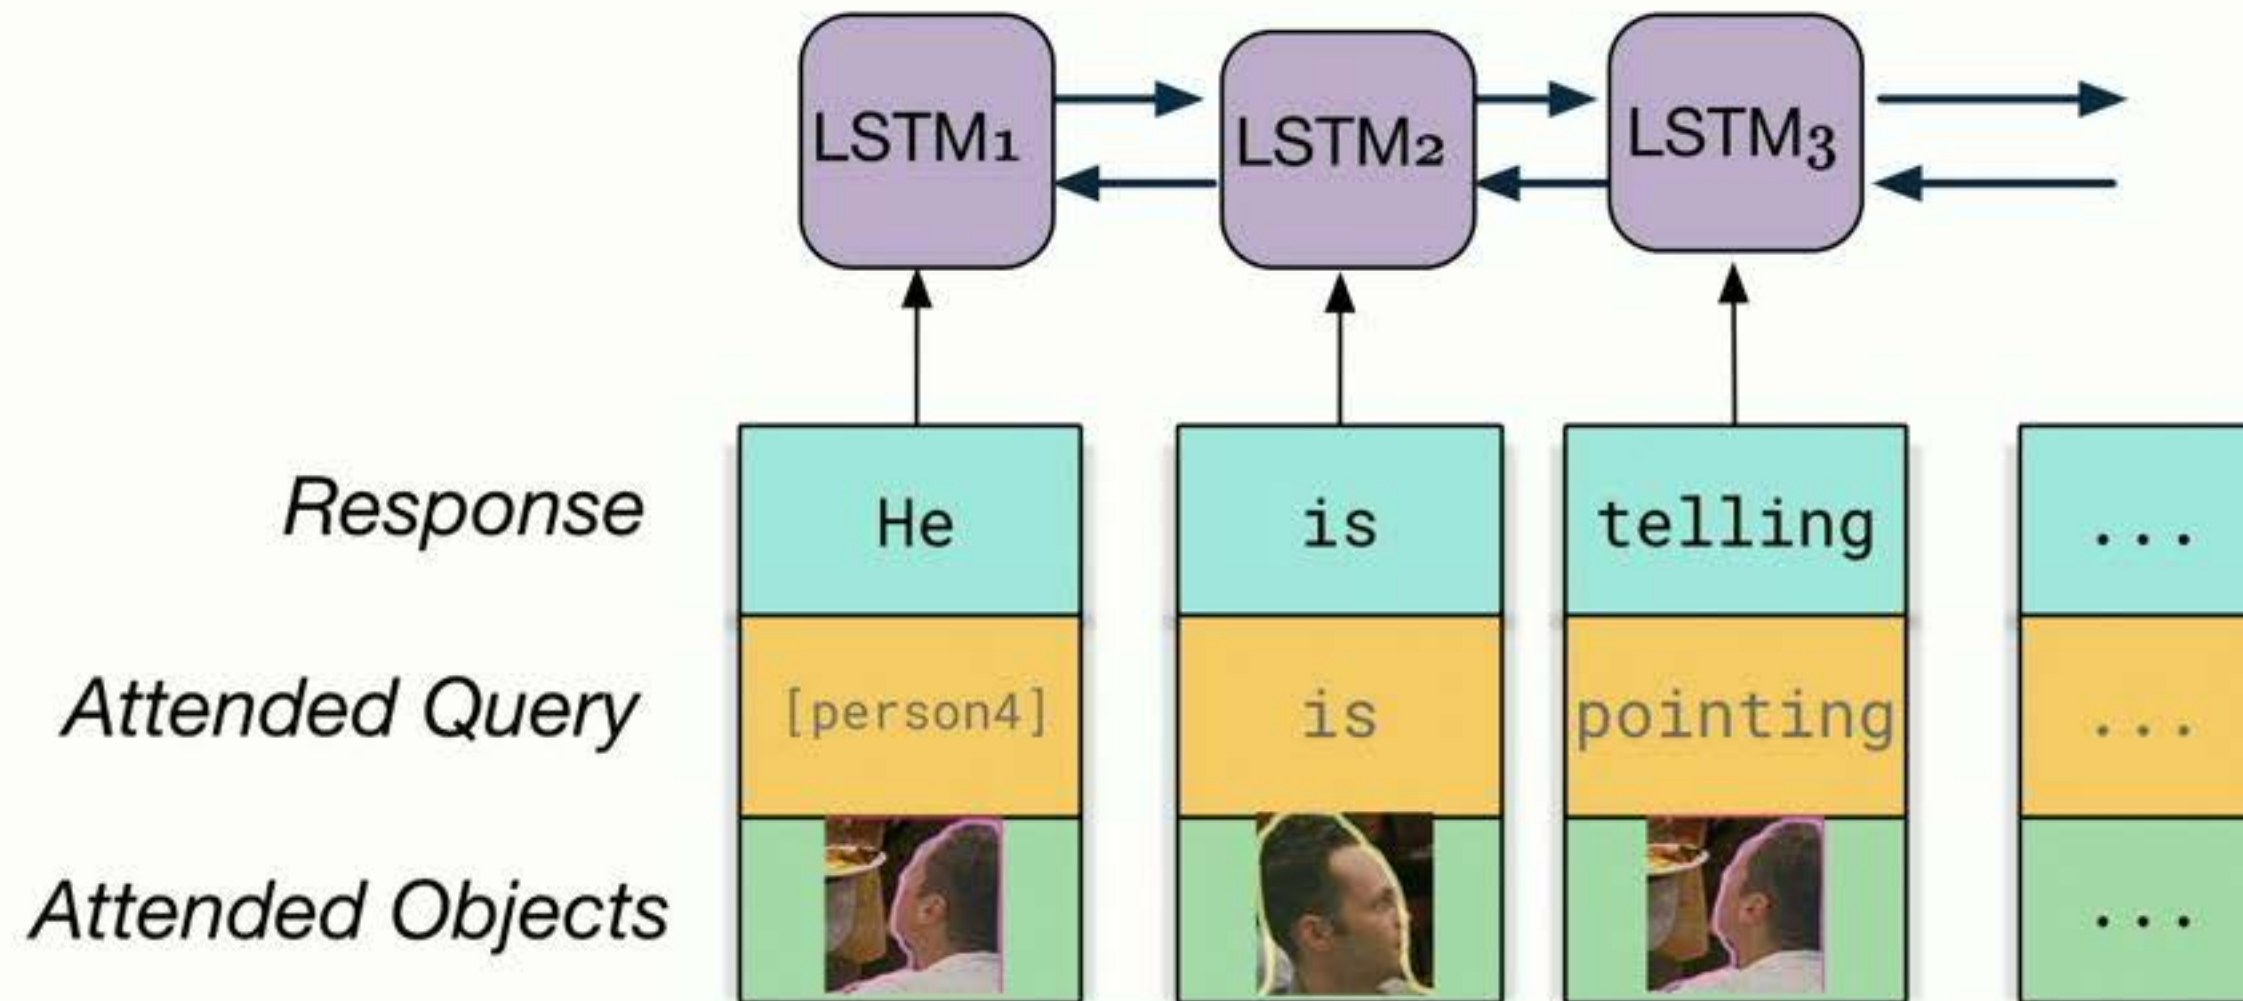
He	is	telling	...
[person4]	is	pointing	...
			...

Response Choice


 is telling **[person3]** 
 that **[person1]** 
 ordered pancakes.



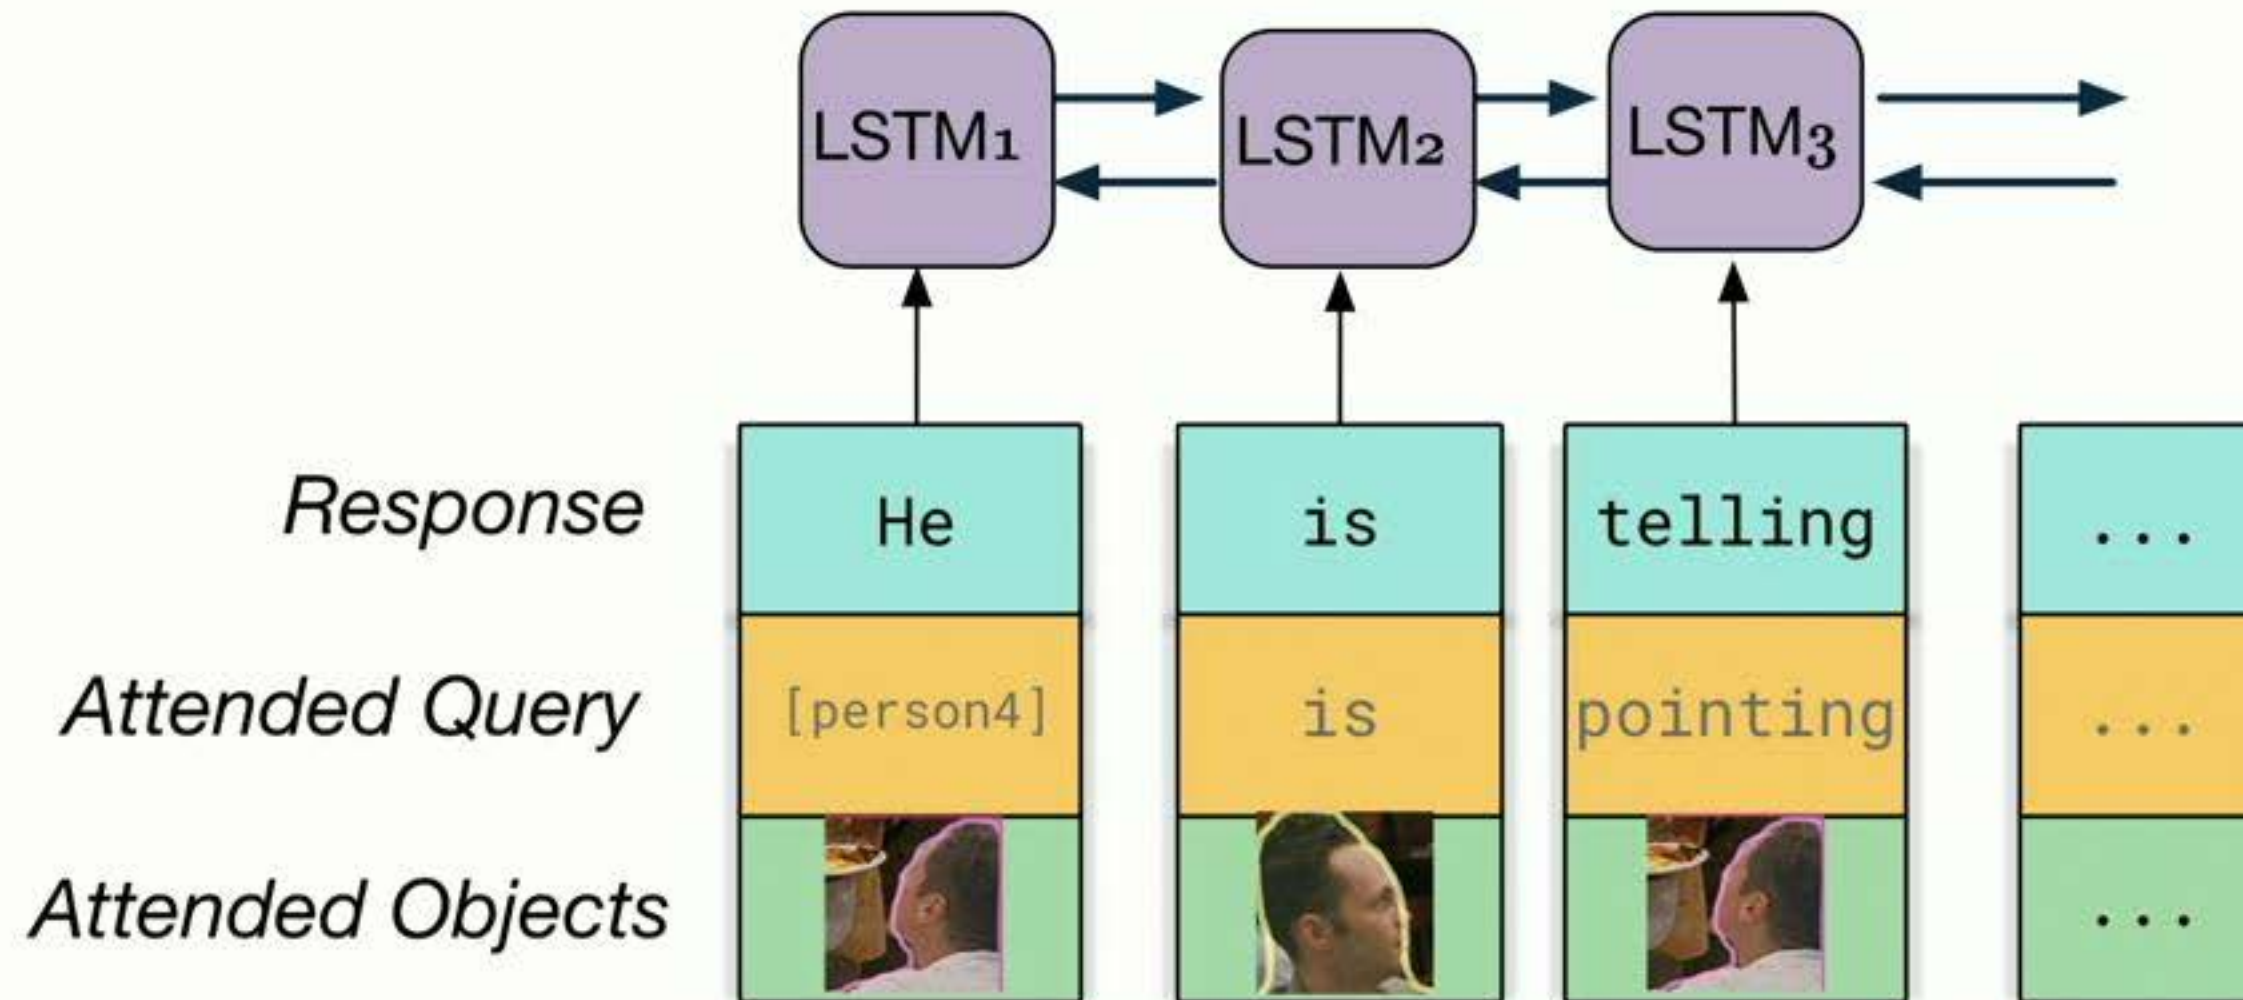
Part 3: Reasoning






Part 3: Reasoning

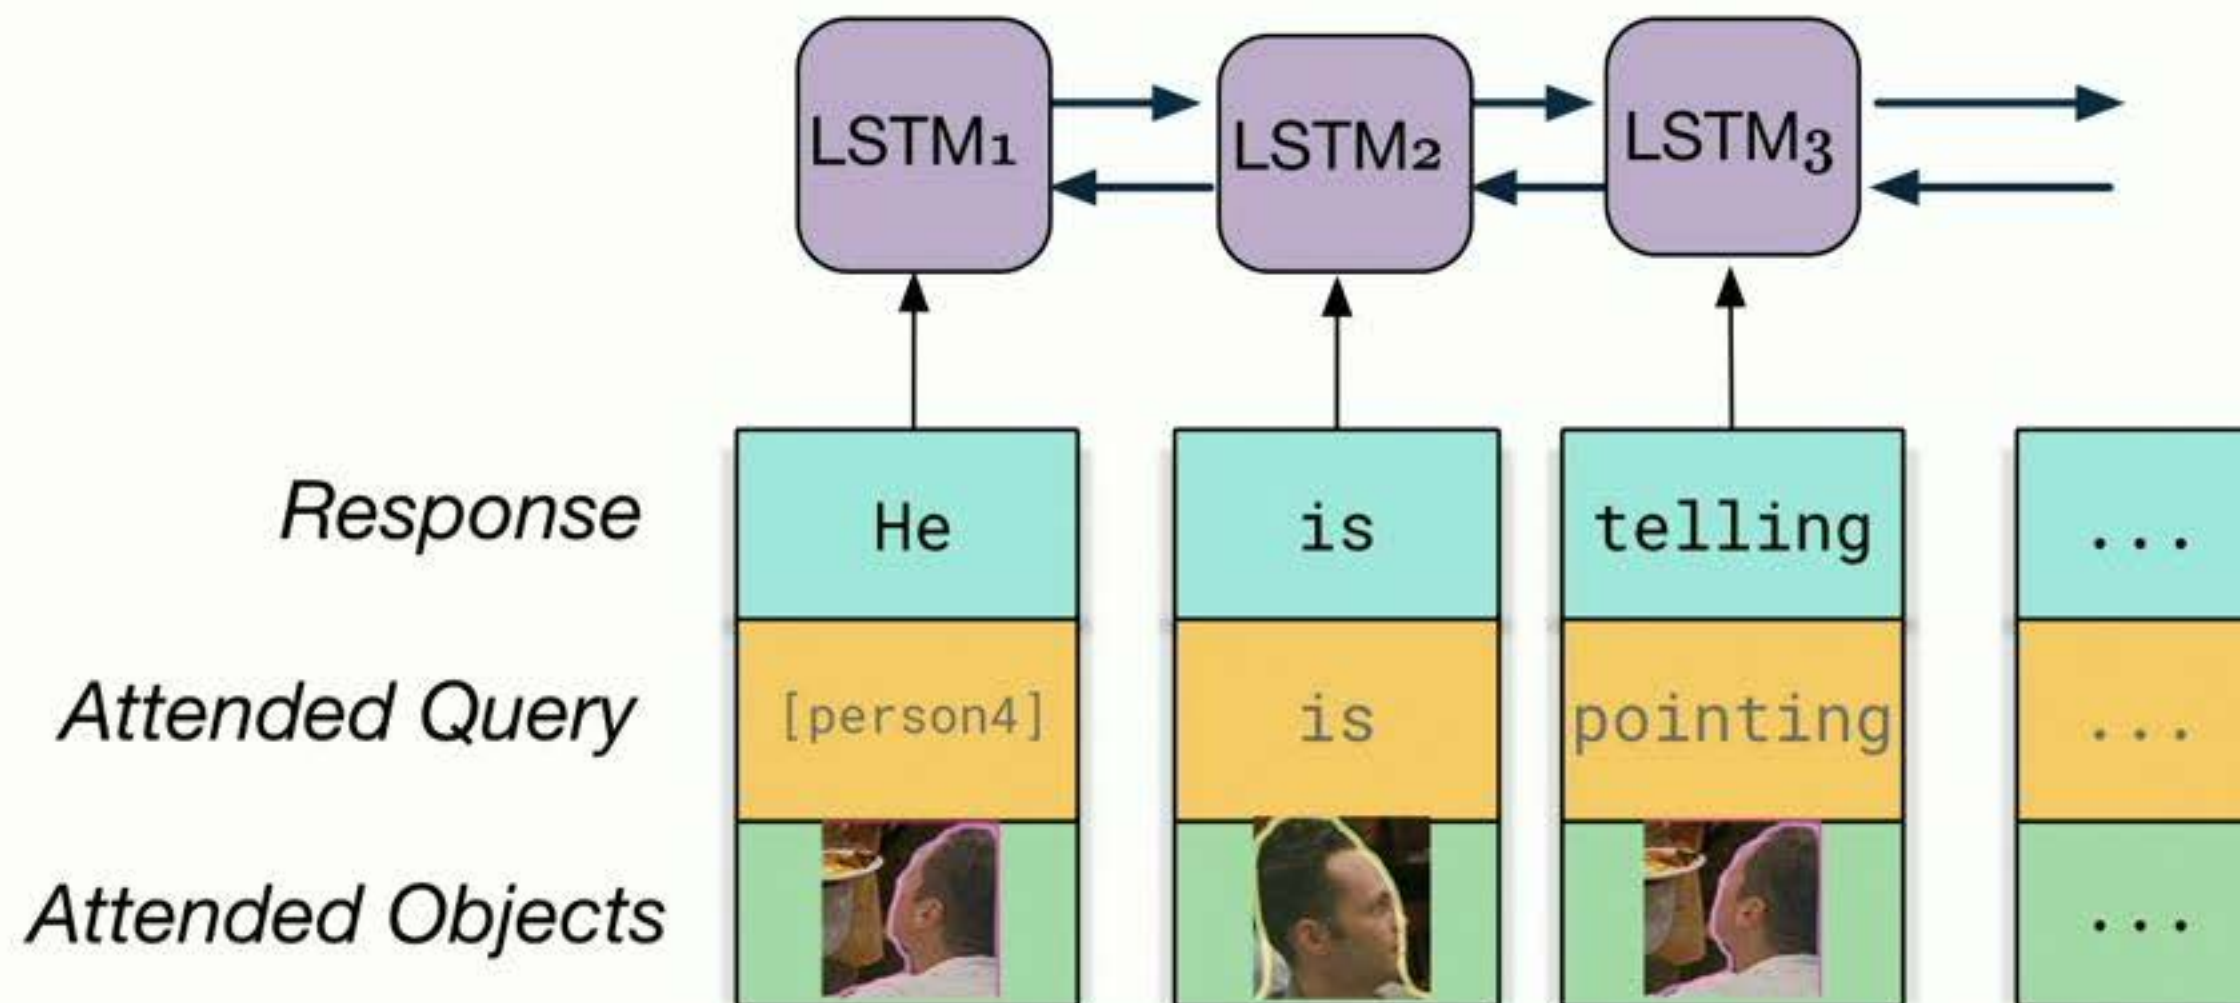
Max pool+multilayer perceptron





Part 3: Reasoning

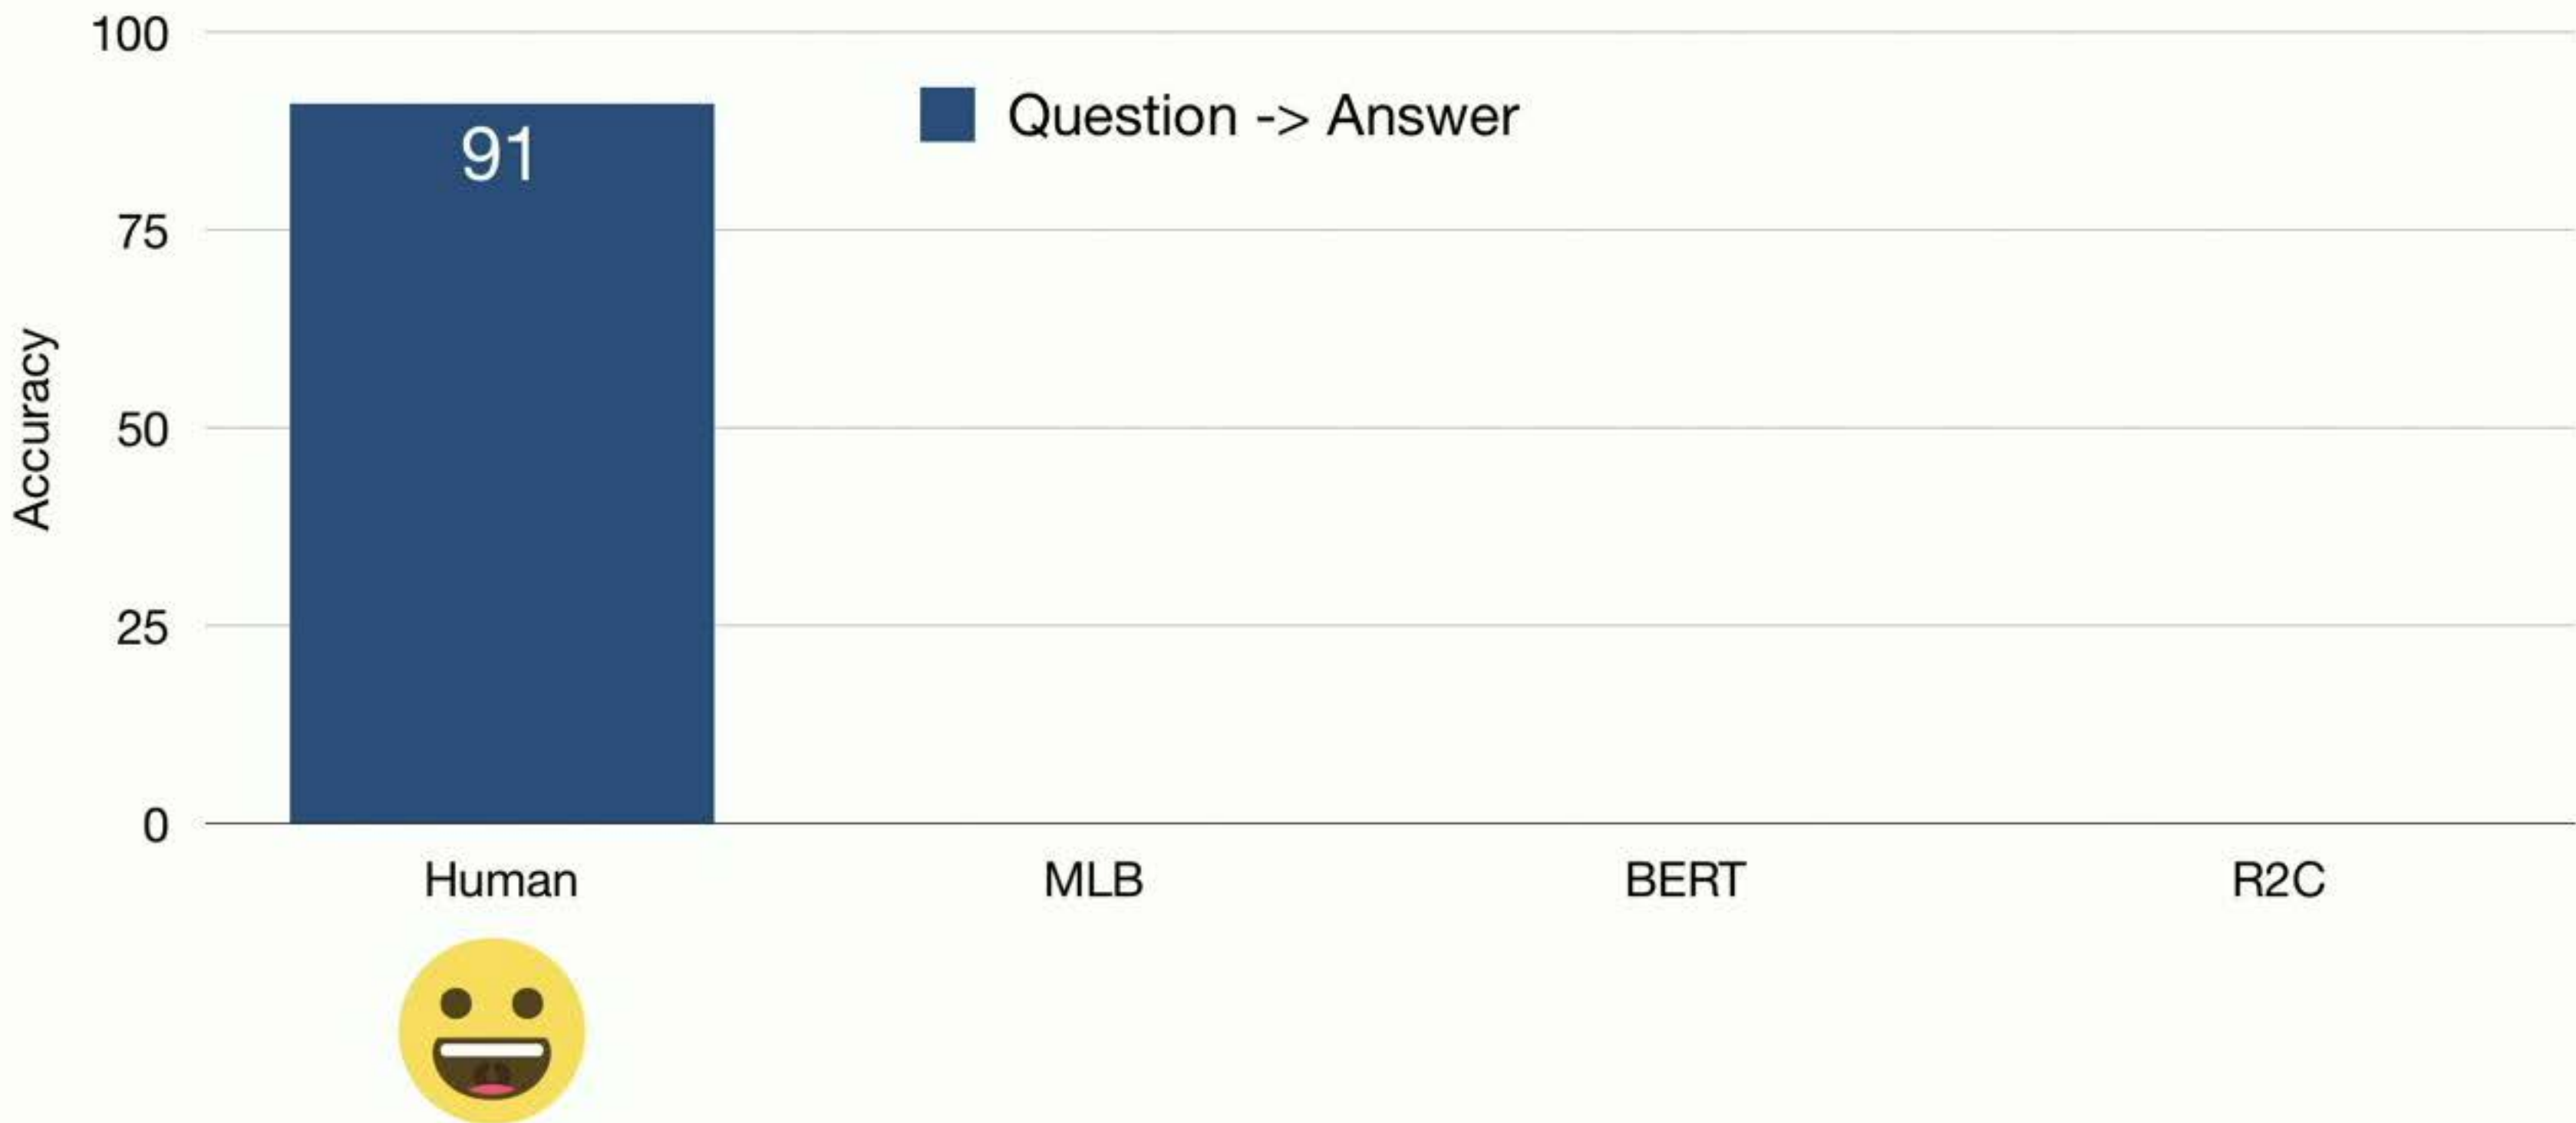
Max pool+multilayer perceptron



VCR Results




VCR Results




VCR Results




VCR Results




VCR Results



VCR Results



VCR Results




[person2]

[tie2]

[person1]

[person3]



[person2]

[person1]

[person3]

[tie2]

1. Why is [person1] pointing a gun at [person2] ?

a) [person1] wants to kill [person2] . 1.4%

b) [person1] and [person3] are robbing the bank and [person2] is the bank manager. 71.7%

c) [person2] has done something to upset [person1] . 18.7%

d) Because [person2] is [person1] 's daughter. [person1] wants to protect [person2] . 8.2%




1. Why is [person1] pointing a gun at [person2] ?

- a) [person1] wants to kill [person2] . 1.4%
- b) [person1] and [person3] are robbing the bank and [person2] is the bank manager. 71.7%
- c) [person2] has done something to upset [person1] . 18.7%
- d) Because [person2] is [person1] 's daughter. [person1] wants to protect [person2] . 8.2%

- a) [person1] is chasing [person1] and [person3] because they just robbed a bank. 33.8%
- b) Robbers will sometimes hold their gun in the air to get everyone's attention. 5.3%
- c) The vault in the background is similar to a bank vault. [person3] is waiting by the vault for someone to open it. 49.1%
- d) A room with barred windows and a counter usually resembles a bank. 11.7%

visualcommonsense.com



HomePaperDownloadCodeExploreLeaderboard

VCR Exploration

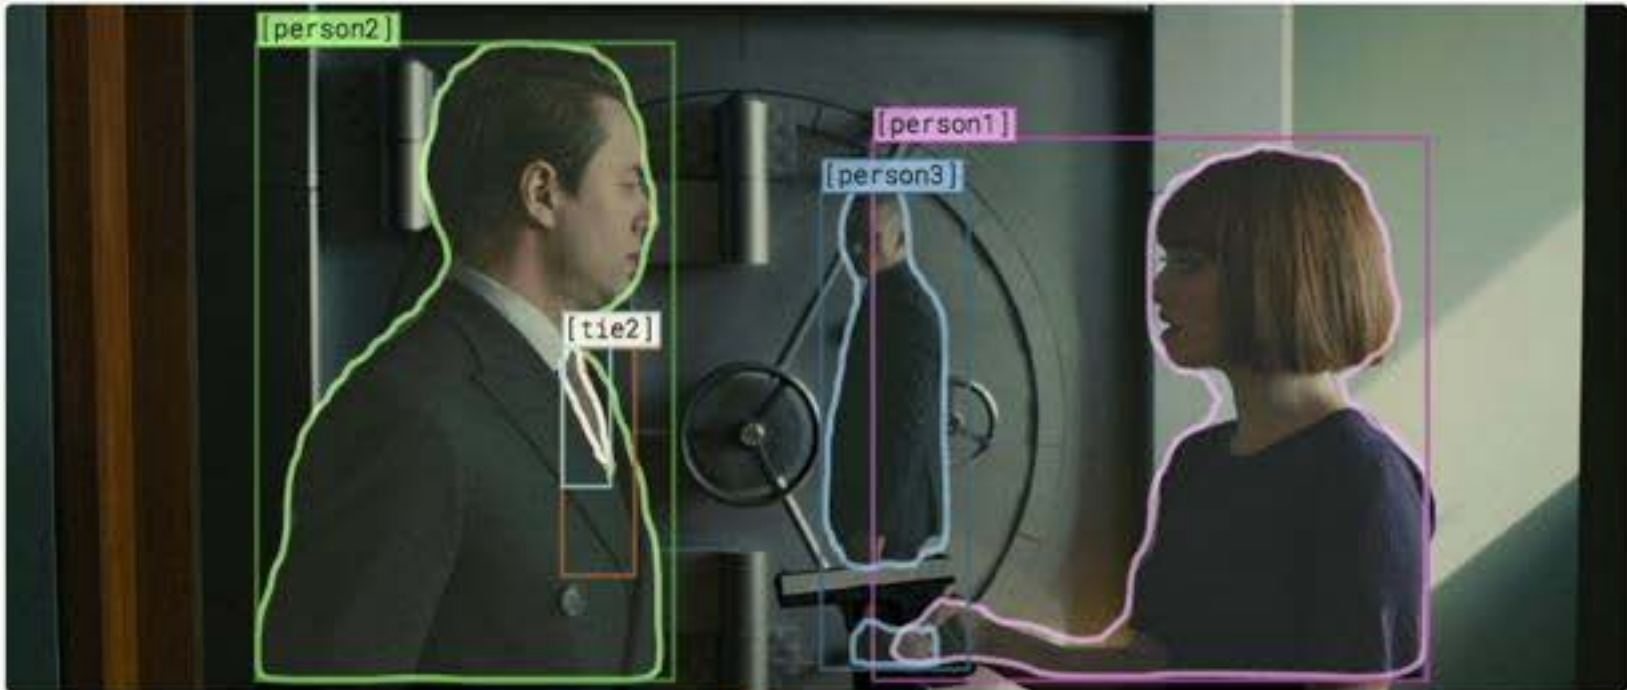
Use the UI below to browse the predictions of our model, R2C, on the validation set. There are 9929 images in the validation set to explore.

For examples from the paper, see [3880 \(example a\)](#), [2518 \(example b\)](#), [6310 \(example c\)](#), or [8512 \(example d\)](#).

⏮ ⏪ ⏩ ⏭

3880

Go ▶



hide allshow all

[person1][person2][person3][tie1][tie2]

1. Why is [person1] pointing a gun at [person2] ?

a) [person1] wants to kill [person2] .	1.4%
b) [person1] and [person3] are robbing the bank and [person2] is the bank manager.	71.7%
c) [person2] has done something to upset [person1] .	18.7%
d) Because [person2] is [person1] 's daughter, [person1] wants to protect [person2] .	8.2%

I think so because...

a) [person1] is chasing [person1] and [person3] because they just robbed a bank.	33.8%
b) Robbers will sometimes hold their gun in the air to get everyone's attention.	5.3%
c) The vault in the background is similar to a bank vault. [person3] is waiting by the vault for someone to open it.	49.1%
d) A room with barred windows and a counter usually resembles a bank.	11.7%

visualcommonsense.com



Home

Paper

Download

Code

Explore

Leaderboard

Visual Commonsense Reasoning

On VCR, a model must not only answer commonsense visual questions, but also provide a rationale that explains why the answer is true.

[Read our paper on arXiv»](#)

Submitting to the leaderboard

Submission is easy! You just need to email [Rowan](#) with your predictions. Formatting instructions are below:

[Submit to the leaderboard»](#)

What kinds of submissions are allowed?

The only constraint is that your system must predict the **answer first**, then the rationale. (The rationales were selected to be highly relevant to the correct Q,A pair, so they leak information about the correct answer.)

- To deter this, the submission format involves submitting predictions for each possible rationale, conditioned on each possible answer.
- A simple way of setting up the experiments (used in the paper) is to consider a task with *query* and four *response* choices. For Q->A the query is the question, and the response choices are the answers. For QA->R, the query is the question and answer, concatenated together, and the response choices are the rationales.

Questions?

If it's not about something private, check out the google group below:


[Get help via the google group»](#)

VCR Leaderboard

There are two different subtasks to VCR:

- **Question Answering (Q->A):** In this setup, a model is provided a question, and has to pick the best answer out of four choices. Only one of the four is correct.
- **Answer Justification (QA->R):** In this setup, a model is provided a question along with the correct answer, and it has to justify it by picking the best rationale out of four choices.

We combine the two parts with the **Q->AR** metric, in which a model only gets a question right if it answers correctly *and* picks the right rationale. Models are evaluated in terms of accuracy (%). How well will your model do?

Rank	Model	Q->A	QA->R	Q->AR
	Human Performance <i>University of Washington</i> (Zellers et al. '18)	91.0	93.0	85.0
1	 CKRE <i>Peking University</i> Feb 9, 2019	66.1	68.5	45.5
2	Recognition to Cognition Networks <i>University of Washington</i> https://github.com/rowanz/r2c Nov 28, 2018	65.1	67.3	44.0
3	BERT-Base <i>Google AI Language</i> (experiment by Rowan) https://github.com/google-research/bert Nov 28, 2018	53.9	64.5	35.0
4	MLB <i>Seoul National University</i> (experiment by Rowan) https://github.com/inhwkim Nov 28, 2018	46.2	36.8	17.2

visualcommonsense.com



Home

Paper

Download

Code

Explore

Leaderboard

Visual Commonsense Reasoning

On VCR, a model must not only answer commonsense visual questions, but also provide a rationale that explains why the answer is true.

[Read our paper on arXiv»](#)

Submitting to the leaderboard

Submission is easy! You just need to email [Rowan](#) with your predictions. Formatting instructions are below:

[Submit to the leaderboard»](#)

What kinds of submissions are allowed?

The only constraint is that your system must predict the **answer first**, then the rationale. (The rationales were selected to be highly relevant to the correct Q,A pair, so they leak information about the correct answer.)

- To deter this, the submission format involves submitting predictions for each possible rationale, conditioned on each possible answer.
- A simple way of setting up the experiments (used in the paper) is to consider a task with *query* and four *response* choices. For Q->A the query is the question, and the response choices are the answers. For QA->R, the query is the question and answer, concatenated together, and the response choices are the rationales.

Questions?

If it's not about something private, check out the google group below:


[Get help via the google group»](#)

VCR Leaderboard

There are two different subtasks to VCR:

- **Question Answering (Q->A):** In this setup, a model is provided a question, and has to pick the best answer out of four choices. Only one of the four is correct.
- **Answer Justification (QA->R):** In this setup, a model is provided a question along with the correct answer, and it has to justify it by picking the best rationale out of four choices.

We combine the two parts with the **Q->AR** metric, in which a model only gets a question right if it answers correctly *and* picks the right rationale. Models are evaluated in terms of accuracy (%). How well will your model do?

Rank	Model	Q->A	QA->R	Q->AR
	Human Performance <i>University of Washington</i>	91.0	93.0	85.0
1	 CKRE <i>Peking University</i> Feb 9, 2019	66.1	68.5	45.5
2	Recognition to Cognition Networks <i>University of Washington</i> https://github.com/rowanz/r2c Nov 28, 2018	65.1	67.3	44.0
3	BERT-Base <i>Google AI Language</i> (experiment by Rowan) https://github.com/google-research/bert Nov 28, 2018	53.9	64.5	35.0
4	MLB <i>Seoul National University</i> (experiment by Rowan) https://github.com/inhwkim Nov 28, 2018	46.2	36.8	17.2

What's next?

Future work

- VCR is a new testbed for commonsense visual reasoning!
- What kinds of new models will do well?



Summary

SWAG



Summary

SWAG





Thanks all!!

More at rowanzellers.com

twitter: @rown




Yonatan Bisk



Roy Schwartz





Why is [person4 ] pointing at [person1 ]?

VCR Results

